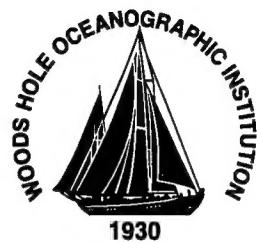


**Massachusetts Institute of Technology
Woods Hole Oceanographic Institution**



**Joint Program
in Oceanography/
Applied Ocean
Science
and Engineering**



DOCTORAL DISSERTATION

Free-Surface Turbulence and Air-Water Gas Exchange

by

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Sean Patrick McKenna

September 2000

20010702 042

MIT/WHOI

00-29

Free-Surface Turbulence and Air-Water Gas Exchange

by

Sean Patrick McKenna

**Massachusetts Institute of Technology
Cambridge, Massachusetts 02139**

and

**Woods Hole Oceanographic Institution
Woods Hole, Massachusetts 02543**

September 2000

DOCTORAL DISSERTATION

Funding was provided by the Woods Hole Oceanographic Institution Ocean Ventures Fund.

**Reproduction in whole or in part is permitted for any purpose of the United States Government.
This thesis should be cited as: Sean Patrick McKenna, 2000. Free-Surface Turbulence and
Air-Water Gas Exchange. Ph.D. Thesis. MIT/WHOI, 00-29.**

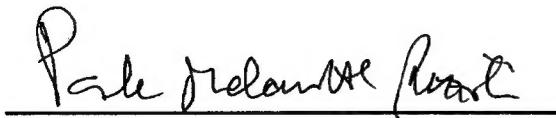
Approved for publication; distribution unlimited.

Approved for Distribution:



Timothy K. Stanton, Chair

Department of Applied Ocean Physics and Engineering



**Paola Malanotte-Rizzoli
MIT Director of Joint Program**



**John W. Farrington
WHOI Dean of Graduate
Studies**

Free-surface turbulence and air-water gas exchange

by

Sean Patrick McKenna

S.M. Oceanographic Engineering, Massachusetts Institute of Technology (1997)

B.S. Mechanical Engineering, Rensselaer Polytechnic Institute (1993)

Submitted to the

Department of Ocean Engineering, MIT

and the

Department of Applied Ocean Physics and Engineering, WHOI

in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN OCEANOGRAPHIC ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

and the

WOODS HOLE OCEANOGRAPHIC INSTITUTION

September 2000

© Sean P. McKenna, 2000

All rights reserved.

The author hereby grants to MIT and WHOI permission to reproduce and
distribute publicly paper and electronic copies of this thesis in whole or in part.

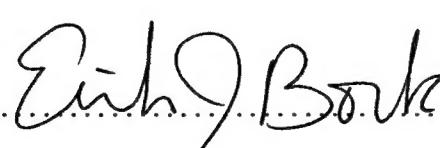
Author 

Department of Applied Ocean Physics and Engineering, WHOI
Department of Ocean Engineering, MIT

August 18, 2000

Certified by 

Dr. Wade R. McGillis
Associate Scientist, WHOI
Thesis Supervisor

Certified by 

Dr. Erik J. Bock
Scientist, Universität Heidelberg
Thesis Supervisor

Accepted by 

Prof. Michael S. Triantafyllou
Chairman, Joint Committee for Applied Ocean Science and Engineering
Massachusetts Institute of Technology/Woods Hole Oceanographic Institution

Free-surface turbulence and air-water gas exchange

by

Sean Patrick McKenna

Submitted to the Department of Applied Ocean Physics and Engineering, WHOI
and the

Department of Ocean Engineering, MIT
on August 18, 2000, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Oceanographic Engineering

Abstract

This thesis investigates the physical mechanisms of air-water gas transfer through direct measurements of turbulence at the air-water interface. To enable this study, a new approach to the particle image velocimetry (PIV) technique is developed in order to quantify free-surface flows. Two aspects of this work are innovative. First, the use of a three-dimensional laser light cone and optical filtering of the camera allow for the motion of fluorescent flow tracers at the water surface to be isolated and measured. Validation experiments indicate that this measurement reflects the fluid motion within the upper few hundred microns. A key benefit to this approach is the ability to deal with deforming surfaces, provided the amplitudes are not prohibitively large. This feature was used in this thesis to explore the surface flow induced by mechanically generated waves. Second, a new hybrid PIV image processing algorithm was developed that provides high accuracy velocity estimation with improved computational efficiency. This algorithm combines the concepts of dynamic Fourier-domain cross-correlation with a localized direct multiplicative correlation.

In order to explore relationships between free-surface hydrodynamics and air-water gas transfer, an oscillating grid-stirred tank was constructed. By its design, this tank can be managed for chemical cleanliness, offers an unobstructed free surface, and is suited for turbulent mixing and air-water gas-exchange studies. A series of acoustic Doppler velocimeter, PIV, and infrared imaging experiments are presented that characterize the flow for the grid forcing conditions studied. Results indicate that the flows are stationary and reasonably repeatable. In addition, the flows exhibit near-isotropic turbulence and are quasi-homogeneous in horizontal planes. Secondary circulations are revealed and investigated.

Finally, PIV measurements of free-surface turbulence are performed with concurrent measurements of gas transfer in the grid tank for a range of turbulent mixing and surface conditions. Surface turbulence, vorticity, and divergence are all affected by the presence of a surface film, with significant effects realized for relatively small surface pressures. Results show that while a relationship between surface turbulence and the gas-transfer velocity is an obvious improvement over that found using an estimate of the bulk flow turbulence, this relationship is dependent on the flow regime. This is revealed through additional surface wave studies. However, the data from both the wave experiments and the grid turbulence experiments can be reconciled by a single relationship between the gas-transfer

velocity and the 1/2-power of the surface divergence, which agrees with previous conceptual models. These results (1) further our understanding of interfacial transport processes, (2) demonstrate the important role of surface divergence in air–water gas exchange, and (3) relate, in a physically meaningful way, the interactions between surface renewal, surfactants, and gas transfer.

Thesis Supervisor: Dr. Wade R. McGillis
Title: Associate Scientist, WHOI

Thesis Supervisor: Dr. Erik J. Bock
Title: Scientist, Universität Heidelberg

For my patient and loving wife

Acknowledgments

During the early course of some of this thesis work, I was supported as an Office of Naval Research Graduate Fellow. This assistance was the impetus for my pursuit of a doctoral degree and is gratefully acknowledged. Very special thanks also to the WHOI Ocean Ventures Fund Program, Mr. F. Thomas Westcott, and the WHOI Education Department for generous financial support over the past several years.

I am grateful to my advisors, Dr. Erik Bock and Dr. Wade McGillis. In addition to giving me my start in the WHOI/MIT Joint Program, both have extended me unwavering support these last six years and have taught me a great deal about innovative experimental science. I also would like to extend my gratitude to the other members of my thesis committee, Dr. James Ledwell, Dr. Eugene Terray, and Prof. Dick Yue. Their comments and suggestions improved my work and this thesis.

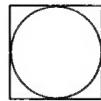
There are a number of people who provided me with assistance during this thesis work, and I am grateful to them all. Steve Fries for his assistance with some of the CRL experiments and subsequent helpful discussions; Jay Sisson for guiding me with the CRL LDV system and racetrack flume; Chris Zappa for very recent and extremely fruitful discussions; Stace Beaulieu for thoroughly reading several chapters of this thesis and providing very useful feedback; Prof. Jerome Milgram and Nicole Suoja for the transport and use of their lab's linear wave tank. I also would like to thank Elliot Rittenberg of IRcameras.com who generously loaned me an infrared imager and video capture package for a long weekend so that I could collect some very interesting free-surface thermal imagery. I am also indebted to Dr. Nelson Frew and Bob Nelson for their assistance with the surface chemistry aspects of this thesis.

I thank my labmates—Korn, Staind, and Incubus—for always keeping me company, day and night, in Smith 117. Thanks also to the Smith security team and the Bigelow and Smith evening crews. It was always nice to see someone else after hours and be greeted with a smile and a hello—that was much appreciated.

Probably the most lasting impact of this thesis experience will be the valuable friendships that I have made in the Joint Program. Kyle Becker, Steve Jayne, and Joe Warren have been great friends and have provided many sources of much appreciated extracurricular activity.

Jason Gobat also has been an invaluable friend as well as an excellent resource over the years. And at MIT, I always have been able to count on the friendship and assistance of Alex Techet. I value these friendships, and all the others I have made these past six years, a great deal.

More than anything, this thesis has taught me a lot about myself—who I am and who I am not. I am hopeful that this journey in self-discovery will help me in my future endeavors. This experience also has shown me how fortunate I am to have been blessed with such a wonderful family. From the beginning, my mother has been my greatest supporter, and her love, encouragement, and understanding have been essential. My grandmother, too, has always been a loving source of quiet strength. I am also very thankful for the support and love of Rich and Ginger Rowcroft. Lastly, I thank my wife, Jessica, for her immeasurable support, patience, and love. This thesis has been difficult not only for me, but also for her, and I share this accomplishment with her.



Contents

1	Introduction	23
1.1	Context and motivation	23
1.2	Organization of the thesis	27
2	Gas transport across air–water interfaces	29
2.1	Introduction	29
2.2	Interfacial transport	30
2.2.1	A simple physical description of air–water transport	30
2.2.2	Free-surface boundary layers	34
2.3	Surface film effects	38
2.3.1	Surface tension	38
2.3.2	Surfactants	40
2.3.3	Marangoni forces and surface elasticity	42
2.4	Models of air–water gas transport	45
2.4.1	Conceptual models	46
2.4.2	Eddy diffusivity models	47
2.4.3	Hydrodynamic models	49
2.4.4	Other parameterizations	50
2.4.5	A model of gas transport including surface film effects	51
2.4.6	Summary discussion	53
3	Performance of digital image velocimetry processing techniques	55
3.1	Introduction	55

3.2	Processing digital image data for velocimetry	56
3.2.1	Cross-correlation algorithms for DPIV	57
3.2.2	Particle tracking	62
3.3	Accuracy and efficiency	64
3.3.1	Uniform translation	66
3.3.2	Uniform rotation	75
3.3.3	Divergent flows	84
3.4	Summary	87
4	A DPIV technique for free-surface flow measurements	89
4.1	Introduction	89
4.2	Description of the surface PIV technique	92
4.2.1	Laser illumination	92
4.2.2	Seed particles	93
4.2.3	Optics	94
4.3	Theoretical context	96
4.3.1	Reflection at an air–water interface	96
4.3.2	Light attenuation in water	97
4.3.3	Scattering of light by particles	98
4.4	Validation of the technique: Phase 1	103
4.4.1	Experimental facility	104
4.4.2	Laser Doppler velocimeter	105
4.4.3	Digital particle image velocimeter	105
4.4.4	Experimental arrangement and procedure	106
4.4.5	Analysis and results	108
4.5	Validation of the technique: Phase 2	114
4.5.1	Experimental arrangement and procedure	115
4.5.2	Analysis and results	116
4.6	Validation of the technique: Phase 3	119
4.6.1	PIV–PIV comparison	119
4.6.2	LDV–PIV comparison	125

4.7	Summary	127
5	Free-surface PIV: application to plane progressive water waves	129
5.1	Introduction	129
5.2	Experimental arrangement and procedure	129
5.2.1	Wave tank and wavemaker	129
5.2.2	Surface PIV setup	131
5.2.3	Wave slope measurement	131
5.3	Results and analysis	133
5.3.1	Wave frequency	135
5.3.2	Wavelength	135
5.3.3	Wave surface velocity field	137
5.3.4	Sources of error	139
5.4	Summary	141
6	Oscillating grid-stirred turbulence	143
6.1	Introduction	143
6.2	Grid tank design and fabrication	145
6.2.1	Tank and grid design	145
6.2.2	Grid drive design	149
6.2.3	Grid forcing conditions and empirical turbulence estimates	152
6.2.4	Tank cleaning and preparation	154
6.3	ADV measurements of bulk flow characteristics	155
6.3.1	The ADV sensor	156
6.3.2	ADV sampling error estimate	158
6.3.3	Velocity time-record requirements	158
6.3.4	Flow repeatability	162
6.3.5	Stationarity of the flow	164
6.3.6	Spatial variability of the flow	166
6.3.7	Discussion	178
6.4	PIV measurements of bulk flow structure	179

6.4.1	Experimental arrangement and procedure	180
6.4.2	Bulk PIV results for $Re = 596$	182
6.4.3	Bulk PIV results for $Re = 730$	191
6.4.4	Discussion	197
6.5	Infrared imagery of free-surface features	198
6.5.1	The infrared imager and experimental arrangement	199
6.5.2	Results and analysis	201
6.6	Summary	212
7	Free-surface hydrodynamics and air–water gas transfer	215
7.1	Introduction	215
7.2	Earlier efforts and motivation for the present approach	215
7.3	Characteristics of the water surfaces studied	222
7.3.1	Water selection	223
7.3.2	Surface tension and surface pressure measurements	223
7.3.3	Surface properties of natural water surfaces	230
7.4	Grid-stirred free-surface turbulence	231
7.4.1	Experimental approach	233
7.4.2	Surface conditions studied	236
7.4.3	Surface hydrodynamics for cleaned and surfactant surfaces	237
7.4.4	Flow fields for cleaned and surfactant surfaces	244
7.5	Measurement of gas transfer in the grid tank	248
7.5.1	The transport model used for gas-transfer measurements	248
7.5.2	The dissolved oxygen sensor	250
7.5.3	Gas-transfer measurement procedures and diagnostics	251
7.6	Grid turbulence gas-transfer rates for cleaned conditions	256
7.7	Effect of PIV particle seeding on gas-transfer velocity	259
7.8	Grid-stirred turbulence and air–water gas transfer	261
7.8.1	Typical experimental procedure	261
7.8.2	Relationships between flow dynamics and gas-transfer velocity	263
7.8.3	Grid tank surface waves and air–water gas transfer	267

7.8.4	Effect of surface films on gas transfer	273
7.8.5	Relationship between surface divergence and gas-transfer	281
7.9	Summary	283
8	Concluding discussion	285
8.1	Summary of key results	285
8.2	Limitations and thoughts for future work	288
A	Relating Re_ℓ and Re_λ	291
B	Uncertainty analysis for chapter 7 PIV results	293
B.1	Velocity uncertainty	293
B.2	Enstrophy uncertainty	295
B.3	Divergence uncertainty	297
References		299

List of Figures

1-1	Influence of surface films of air–water gas transfer	25
2-1	Sketch of air–water gas transport process	31
2-2	Diagram of the Wilhelmy plate technique	39
2-3	Molecule of oleic acid and preferred insoluble surfactant orientation	41
2-4	Conceptualization of surface renewal by a turbulent eddy in the presence of a surface film	42
3-1	Schematic of direct spatial cross-correlation calculation	58
3-2	Variation of RMS error and mean bias error for all four DPIV methods over the 0–1 px displacement range with 32×32 subimage processing	67
3-3	Displacement frequency distributions for FFT and direct spatial correlation methods	70
3-4	Variation of RMS error and mean bias error for all four DPIV methods over the 0–10.25 px displacement range with 32×32 subimage processing	72
3-5	Flume calibration experiment result	73
3-6	Variation of relative error over the 0–1 px displacement range and the 0–5.25 px displacement range for all four DPIV methods with 16×16 subimage processing	74
3-7	Variation of relative error over the 0–1 px displacement range and the 0–10.25 px displacement range for all four DPIV methods with 32×32 subimage processing	75

3-8 Variation of relative error over the 0–1 px displacement range and the 0–10.25 px displacement range for all four DPIV methods with 64×64 subimage processing	76
3-9 Variation of RMS error, mean bias error, and relative error for all four DPIV processing methods over the 5–6 px displacement range with 32×32 subimages	77
3-10 Same as figure 3-2 but with raw, ungridded PTV results included	79
3-11 Variation of relative error over the 0–1 px displacement range, the 0–10.25 px displacement range, and the 5–6 px displacement range, showing raw, ungridded PTV performance	80
3-12 Displacement fields for divergent flow A	85
4-1 Basic elements of the (surface) PIV system developed in this thesis	93
4-2 PIV surface flow measurement optical arrangement	94
4-3 Reflection and refraction of light at an air–water interface	96
4-4 Mie scattering diagram for a particle with a diameter of $30 \mu\text{m}$	101
4-5 Dependence of scattered intensity on particle diameter	102
4-6 Experimental setup for the phase 1 comparison between LDV and standard PIV systems studying a channel flow	107
4-7 Horizontal velocity time series showing location of “coincident” samples . .	110
4-8 Comparison between mean velocity components measured by LDV and PIV .	112
4-9 Comparison between turbulent flow statistics measured by LDV and PIV .	113
4-10 Velocity spectra as measured by LDV and PIV in the phase 1 channel flow	114
4-11 Experimental setup for phase 2 comparison between LDV and surface PIV systems studying a channel surface flow	116
4-12 PIV camera image showing camera field of view with LDV beams and measurement volume	117
4-13 Mean surface velocities as measured by LDV, DPIV, and PTV techniques for a basic channel flow	118
4-14 Experimental setup for phase 3 comparison	120

4-15 Sample profiling PIV image and the temporally averaged, mean velocity field result	122
4-16 Result from PIV-PIV surface velocity comparison experiment	123
4-17 Mean surface velocities as measured by LDV, DPIV, and PTV techniques for a wind-driven flow	126
4-18 LDV and PIV time series of wind-driven surface velocity for run no. 4	127
5-1 Diagram of linear wave tank setup	130
5-2 Diagram of wave slope measurement optics	132
5-3 Sample wave surface velocity field	134
5-4 Sample wave surface velocity records computed from the wave slope measurement and the surface PIV velocity measurement	138
6-1 Photograph of the oscillating grid-stirred tank in its final form	147
6-2 Photograph of the final grid drive design	151
6-3 Sketch of the ADV sensor probe head	156
6-4 Locations of the ADV measurements relative to both grid geometries	157
6-5 Temporal behavior of flow statistics from the onset of grid motion measured at location 1 for $Re = 469$	160
6-6 Repeatability measurements for $Re = 555$	163
6-7 Collection of ADV measurements indicating degrees of repeatability	165
6-8 Locally binned flow statistics for $Re = 974$	167
6-9 Turbulent velocity spectra showing the influence of the ADV noise floor	168
6-10 Spatial variability of mean flow components measured at a depth of 8 cm for grid A	170
6-11 Spatial variability of mean flow components measured at a depth of 8 cm for grid A	171
6-12 Spatial variability of turbulent fluctuating velocities measured at a depth of 8 cm for grid A	173
6-13 Spatial variability of turbulent fluctuating velocities measured at a depth of 8 cm for grid A	174

6-14 Spatial variability of mean and turbulent fluctuating components measured at a depth of 8 cm for grid B	175
6-15 Ratios of RMS turbulent velocity fluctuations to mean velocities as a function of grid Reynolds number	176
6-16 RMS turbulent velocity fluctuations as a function of grid Reynolds number	177
6-17 Isotropy of the grid-generated turbulence as a function of grid Reynolds number	178
6-18 Comparison between empirically predicted RMS turbulent velocities and those measured	179
6-19 Horizontal spatial variability of RMS turbulent velocities over all Reynolds number studied	180
6-20 Location of the imaged areas for the bulk flow PIV mapping measurements	181
6-21 Mean flow and turbulent kinetic energy visualization in three vertical planes for $Re = 596$	184
6-22 Vertical decay of the RMS turbulent velocities for $Re = 596$	186
6-23 Overall mean profiles of the RMS horizontal and vertical turbulent fluctuations for $Re = 596$	187
6-24 Vertical dependence of the integral lengthscales for $Re = 596$	189
6-25 Overall mean profiles of the integral lengthscales for $Re = 596$	190
6-26 Mean flow and turbulent kinetic energy visualization in three vertical planes for $Re = 730$	193
6-27 Vertical decay of the RMS turbulent velocities for $Re = 730$	194
6-28 Overall mean profiles of the RMS horizontal and vertical turbulent fluctuations for $Re = 730$	195
6-29 Vertical dependence of the integral lengthscales for $Re = 730$	196
6-30 Overall mean profiles of the integral lengthscales for $Re = 730$	197
6-31 Approximate temperature calibration for the IR imager	202
6-32 Sample IR image of the water surface for $Re = 282$ (grid A)	203
6-33 Persistent thermal features and their variability at the free surface for grid-stirred turbulence	204

6-34 Persistent thermal features and their variability at the free surface for grid-stirred turbulence	205
6-35 Persistent thermal features and their variability at the free surface for grid-stirred turbulence	206
6-36 Persistent thermal features and their variability at the free surface for grid-stirred turbulence	207
6-37 Comparison of surface renewal behaviors for two different grid conditions .	210
7-1 Gas-transfer results from Broecker <i>et al.</i> [12] and Frew <i>et al.</i> [35]	217
7-2 Gas-transfer results from Frew <i>et al.</i> [35] and McGillis [95]	218
7-3 Gas-transfer results from Asher and Pankow [6]	219
7-4 Gas-transfer results from Hara <i>et al.</i> [48] and Saylor and Handler [115] . . .	221
7-5 Surface tension of spring water as a function of temperature	224
7-6 Surface pressure isotherms for laboratory tap water and commercial spring water	225
7-7 Surface pressure isotherms for films of “cholesterol” and oleyl alcohol	226
7-8 Surface pressure isotherms for pure spring water and spring water with cleaned PIV supernatant	228
7-9 Statically derived elasticity isotherms for the two surfactants studied (“cholesterol” and oleyl alcohol) and laboratory tap water	229
7-10 Sample π - A curves for three basic monolayer film types and representative sea-surface microlayer films from the California Bight	232
7-11 Locations of the surface PIV measurement locations	233
7-12 Ratios of the surface RMS fluctuating velocities to the surface mean velocities at the four measurement locations	240
7-13 Horizontal surface RMS fluctuating velocities at the four measurement locations	241
7-14 Isotropy of the surface RMS fluctuating velocities at the four measurement locations	242
7-15 Surface enstrophy at the four measurement locations	243
7-16 Surface divergence at the four measurement locations	245

7-17 Impact of surface pressure on surface hydrodynamics	246
7-18 Surface flow fields for cleaned and surfactant surfaces	247
7-19 Photograph of the water quality sensor used for the bulk measurements of dissolved oxygen and temperature	250
7-20 Diagram of the equilibration system used with the liquid nitrogen gas . . .	252
7-21 Results from the O ₂ sensor relocation experiment	254
7-22 Dependence of the gas-transfer velocity on Reynolds numbers for cleaned conditions	257
7-23 Examples of time-dependent behavior of the gas-transfer velocity	259
7-24 Impact of PIV seed particles on gas-transfer rates	260
7-25 Dependence of the gas-transfer velocity on bulk mixing Reynolds number for grid-stirred turbulence	264
7-26 Dependence of the gas-transfer velocity on RMS surface velocity fluctuations for grid-stirred turbulence	265
7-27 Dependence of the gas-transfer velocity on surface enstrophy for grid-stirred turbulence	266
7-28 Dependence of the gas-transfer velocity on surface divergence for grid-stirred turbulence	267
7-29 Sketch of the wave generation scheme in the grid tank	268
7-30 Dependence of the gas-transfer velocity on RMS surface velocity fluctuations for grid-stirred turbulence and mechanically generated surface waves	270
7-31 Dependence of the gas-transfer velocity on surface enstrophy for grid-stirred turbulence and mechanically generated surface waves	271
7-32 Dependence of the gas-transfer velocity on surface divergence for grid-stirred turbulence and mechanically generated surface waves	272
7-33 Example of a flow field from the grid tank <i>waves</i> data set	273
7-34 Influence of surface pressure on gas-transfer velocity for grid-stirred turbulence and mechanically generated surface waves	274
7-35 Dimensionless relationship between surface divergence and gas-transfer velocity for all data sets	282

List of Tables

3.1	Computational performance of DPIV/PTV processing algorithms	78
3.2	Simulated uniform rotation results	81
3.3	Simulated divergent flow results	86
4.1	Parameters used in theoretical Mie scattering calculations	100
5.1	Wave frequency results	135
5.2	Wave velocity results	139
6.1	Relevant dimensions and parameters for the grid-stirred tank	146
6.2	Grid conditions and characteristic parameters	154
6.3	Results from calm water ADV measurements	159
6.4	RMS errors for mean value estimates determined from a finite record length	161
6.5	Environmental parameters during the IR imagery experiments	208
7.1	Initial surface pressure relative to UHP water for various surfaces studied .	230
7.2	Summary of the surface PIV/gas-transfer experimental conditions	236

Chapter 1

Introduction

1.1 Context and motivation

The interaction between the world ocean and the atmosphere plays a crucial role in many global processes, including life itself. The ocean and the atmosphere are fundamentally a coupled system, and the dynamics of each are closely linked. For example, energy is transferred from the atmosphere to the ocean surface driving the circulation of the upper ocean. In turn, energy from the ocean is fed back to the atmosphere affecting atmospheric circulation, weather and planetary climate (Kraus and Businger [74]). Heat and mass are exchanged continually between the ocean and the atmosphere, and the air-sea exchange of heat, water vapor, and many environmentally relevant gases is fundamental to the climatic balance of the planet. With the current attention to potential trends in the Earth's climate, the field of air-sea interaction has seen growing interest in the past few decades and is expected to continue as an important field of research.

The air-water exchange of gases has wide-ranging environmental implications, from oceanic CO₂ transport and questions of global warming to O₂ transport and the reaeration of lakes and rivers. In light of current worldwide increases in population and industrialization, gas transfer at water surfaces is a critical factor in understanding the transport and fate of various anthropogenic materials in the environment and in their responsible management (*e.g.*, CO₂, sulfur oxides, methane, freons, metals, and the many organic chemicals used in pesticides, herbicides, consumer products, and industrial processes). Additionally, air-

water transport plays an important role in many engineering applications such as separation processes, bubbly flows, and reactors.

In the context of air-sea mass transfer, knowledge of fluxes on large scales has often been the desired goal in order to estimate global budgets. However, it is clearly the air-water interface through which all gaseous, liquid, and particulate material must ultimately pass when exchange between the atmosphere and ocean occurs. In the past, large-scale air-sea interaction studies, such as global climate and ocean circulation models, have relied heavily on the verity of parameterizations of small-scale air-sea exchange processes occurring at the air-water interface. A complete understanding of these complex processes and their parameterization remains elusive, but is essential to address adequately the larger scale ocean-atmosphere problems. Therefore, studies that probe the fundamental interfacial mechanisms of air-water exchange processes are a vital part of improving our understanding of air-sea gas transfer.

For sparingly soluble gases such as CO₂, O₂, and N₂, the transfer across the air-water interface is governed by the physics within a very thin aqueous boundary layer. Because this layer is on the order of a tenth of a millimeter, exploring the underlying mechanisms of air-water transport has been experimentally difficult. Exacerbating this already challenging problem are the effects of surfactants. Surfactants, which can form surface films, can dramatically modify the hydrodynamics of a free surface and now are well-known to have a significant impact on air-water gas transfer. Surfactant sources in the oceans include natural and anthropogenic terrestrial sources, delivered by atmospheric transport and runoff, and marine biological sources. The largest source of surface film material in the ocean is production by marine organisms, most notable phytoplankton, which exude natural surfactants as metabolic by-products. The breakdown of dead organisms also contributes to the surfactant pool (Frew *et al.* [36]). Since such sources are ubiquitous in many regions of the world ocean, so too are surface films. Consequently, cognizance and consideration of the effects of surface films are requisite for comprehensive air-sea interaction studies.

To provide a sense of the challenges associated with studying gas transfer at air-water interfaces, figure 1-1 presents a particularly illustrative result from this thesis. This figure shows results from a turbulence-driven air-water gas-transfer experiment in an oscillating

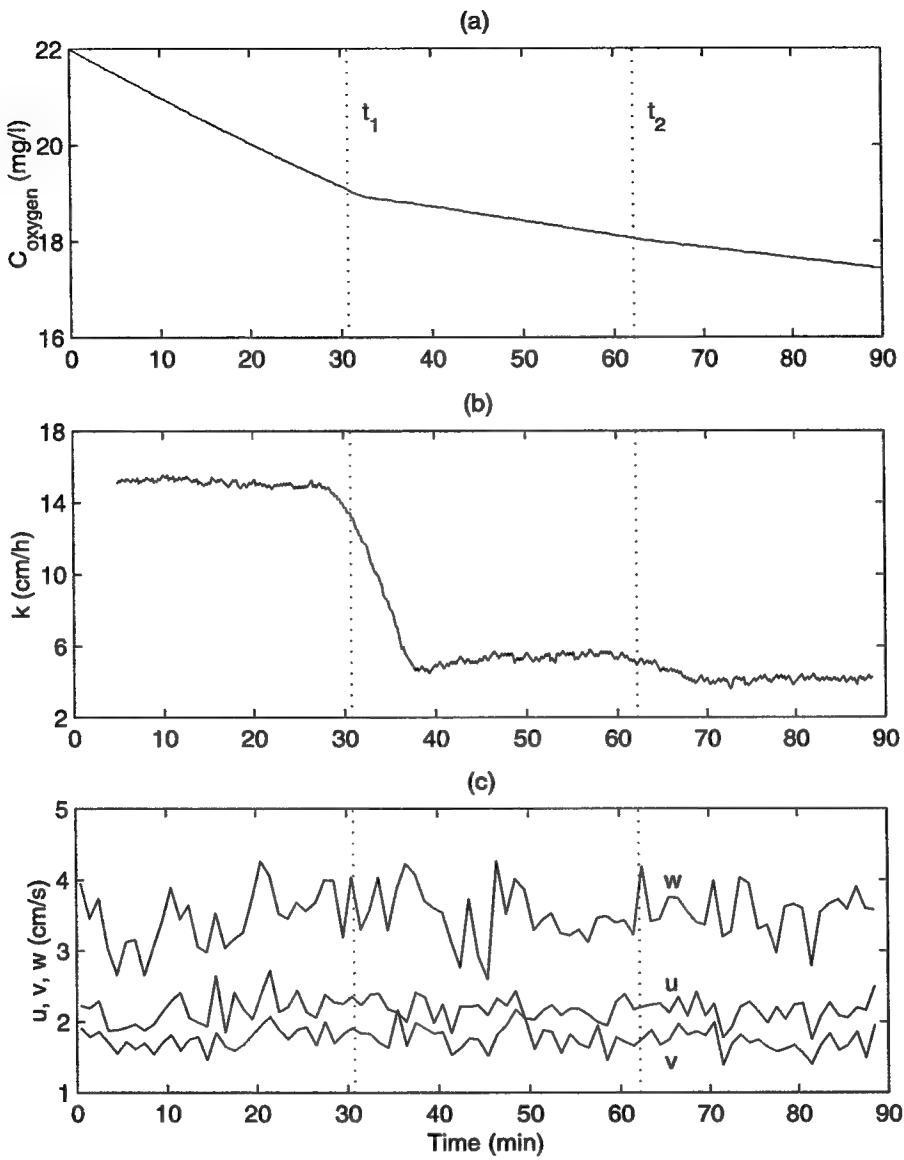


Figure 1-1: Influence of surface films of air-water gas transfer. Results are from an oscillating grid-stirred tank experiment. (a) Concentration of (well-mixed) bulk dissolved oxygen gas, (b) gas-transfer velocity k , and (c) binned root-mean-square turbulent fluctuating velocities (u, v, w) measured at a depth of 16 cm. The time t_1 corresponds to the initial addition of surfactant, and the time t_2 corresponds to the second addition of surfactant. The surfactant added was cholesterol and is discussed in chapter 7.

grid-stirred tank. There are three striking features of these results. First is the clearly marked reduction in the gas-transfer velocity (which is directly related to the gas flux) when a surfactant was added to the water surface. In this case, the reduction was 60%! Second is the amount of surfactant added that produced such a dramatic effect: 340 μg at t_1 , with an estimated surface concentration of 0.165 $\mu\text{g}/\text{cm}^2$ (roughly a 1 mm diameter spherical drop of material over an area of about 2000 cm^2). Lastly, and perhaps most surprisingly, is the complete lack of evidence of a change in the subsurface turbulent velocity record. The velocity measurement was acquired at a single location 16 cm from the free surface, which, as an ocean measurement, would be likely considered “near-surface.” This is somewhat disconcerting: the mass transfer of the system was drastically attenuated, yet the measured system turbulence near the interface remained unchanged. The consequences of this phenomenon for estimating gas transfer from measurements of near-surface turbulence are thought-provoking.

In the field, numerous methods have been used over the years to estimate the gas-transfer velocity from *in situ* data: *e.g.*, monitoring steady-state distributions of natural ^{14}C and bomb- ^{14}C , the ^{222}Rn deficit method, deliberate tracer studies, and more recently, eddy correlation and profiling techniques. The gas-transfer velocity is often parameterized in terms of wind speed (*e.g.*, Wanninkhof [136]) since wind speed is easily measured in the field from ships or buoys, can be estimated from satellite-borne scatterometers, and may be predicted from general atmospheric circulation models (Frew [34]). While convenient, wind speed relationships for the transfer velocity do not account for the spread in current measurements. Both laboratory and field measurements have revealed that a unique relation between wind speed and the transfer velocity is not likely to exist for natural waters containing organic surface-active matter. The reason for this shortcoming, which is not limited to a wind speed relationship, is related to the effect illustrated by figure 1-1.

Fundamentally, the exchange of gases between water and air occurs through thin boundary layers at the air–water interface where turbulent motions are suppressed giving way to molecular diffusion processes. The details of these layers, their thicknesses and the gradients within them, are determined by the aero- and hydrodynamics near the surface. For sparingly soluble gases, the airside effect is less critical, and the transport is determined by

the waterside boundary layer. The fluid flow in this boundary layer is highly sensitive to the presence of a surface film and can be modified considerably, which is the reason gas transfer exhibits a similar sensitivity. Experimental efforts to explore the subsurface flow in this thin region have been thus far challenging. An alternative to subsurface flow measurements are measurements of the free-surface flow itself. At present, such measurements of free-surface flows are still novel, but they are hoped to provide improved insight into the physical mechanisms of air–water gas exchange by probing the very interface where the transport occurs. This thesis develops a means to effectively study free-surface flows with the goal of exploring relationships between air–water gas transfer and free-surface hydrodynamics.

1.2 Organization of the thesis

Eight chapters comprise this thesis. This first chapter introduces the broader context of the focus of this thesis and motivates the present work. Additional background material can be found within individual chapters. Chapter 2 expands on the physics of air–water gas exchange by delving into greater detail. Some of the fundamentals of transport are presented first, followed by an explanation of surfactants and their role in free-surface hydrodynamics and air–water gas transfer. The chapter concludes with a presentation of a number of classical models of interfacial gas transfer.

Chapter 3 presents the analysis of a number of digital image velocimetry processing techniques that are used throughout this thesis. These processing algorithms are part of a digital particle image velocimetry (DPIV) technique that forms the backbone of almost all of the experimental studies in this thesis. The analysis in chapter 3 investigates algorithm accuracy and efficiency for a variety of numerically simulated flow fields.

Chapter 4 is the description and validation of the DPIV technique developed as part of the thesis to study free-surface flows. This technique was developed in order to overcome some of the crucial difficulties in making measurements *at* a free surface. In this chapter, the experimental implementation of the technique is outlined and the principle methodology explained. The remainder of the chapter presents a number of comparative experiments performed to validate the PIV system and its ability to acquire measurements at a free sur-

face. Chapter 5 further extends the analysis and validation of the surface PIV technique by looking at the problem of mechanically generated surface waves. In this chapter, the surface PIV technique is used to measure the characteristics of low amplitude, plane progressive water waves, and the results compared with independent data obtained using a wave slope measurement.

Chapter 6 introduces the experimental tank used to study the relationships between free-surface hydrodynamics, air–water gas transfer, and surfactants. In this chapter, the oscillating grid-stirred tank that was constructed for this thesis is presented and documented. The motivations behind its design are given, and the methods employed to manage the tank for chemical cleanliness are discussed. Next, the results from several diagnostic experiments that examined flow start-up transients, repeatability, and stationarity are reported. In addition, a series of experiments that quantified the turbulence in the grid tank are discussed. DPIV measurements then are shown that explored the flow in several vertical planes of the tank, looking at spatial variability and secondary flows. Finally, infrared imagery of the free surface is shown that provided some insight as to the character of the surface flow field.

Chapter 7 synthesizes the measurements of free-surface hydrodynamics with measurements of air–water gas transfer and surface condition. This chapter shows the basic surface properties of the surfaces studied in the grid-stirred tank experiments. Results from surface PIV measurements of free-surface hydrodynamics for a number of different surface conditions in the grid tank are presented. The procedure for the measurement of gas transfer in the grid tank then is discussed, and some diagnostic results, such as the effect of PIV seed particles on gas transfer, are shown. Finally, the PIV measurements are integrated with the gas-transfer measurements and relationships between the two are explored. In the final chapter, the major findings of this work are summarized. Some possible shortcomings of the thesis are noted, and some thoughts on future work are given.

Chapter 2

Gas transport across air–water interfaces

2.1 Introduction

The transport of gases, and passive scalars in general, across air–water interfaces is important to many fields of science and engineering. As discussed in chapter 1, the exchange of heat and mass between the earth’s atmosphere and its oceans is crucial to the environmental balance of the planet. Similarly, air–water gas exchange plays an important role in the reaeration of lakes and rivers. Many systems used in chemical engineering processes also depend on air–water transport, *e.g.*, bubbly flows, condensers, evaporators, reactors, etc. A detailed understanding of the process of air–water transport therefore is desirable. In this chapter, some fundamental aspects of air–water gas transfer are presented. A simple physical description of the process and the governing equations are discussed first. The subject of surfactants is introduced, and the role of surface films in air–water exchange is described. Several models of the transport process then are presented, and an illustrative model is developed that describes the impact of a surface film on surface renewal and gas transfer.

2.2 Interfacial transport

The transport of gas across an air–water interface ultimately occurs on a molecular level through a process of diffusion. This is true regardless of the nature of the flow regimes on either side of the interface, be they quiescent or turbulent. For the case of sparingly soluble gases, the transfer rate is controlled by the waterside resistance and the airside resistance can be neglected. Under turbulent conditions, the diffusion process results in a very thin aqueous mass boundary layer at the interface, and the concentration gradient across this layer determines the flux of gas across the interface. Waterside mixing processes such as free-surface waves and near-surface turbulence act to thin this boundary layer, steepening the concentration gradient, and enhancing gas exchange.

In this thesis, the case of an unbroken air–water interface is explored. The more involved situation of a water surface that is breaking due to waves or through strong upwelling events, characteristic of many natural systems, is not considered. In such highly dynamic situations, air entrainment, bubble formation, and aerosol generation can become important to the transport process. However, even for bubbles and aerosols, the exchange between air and water is still achieved through diffusion across a two-phase interface. In this chapter, the air–water interface is considered flat, and in the context of air–sea interactions, the underlying flow is taken to be well-mixed.

2.2.1 A simple physical description of air–water transport

To introduce the problem of air–water gas transfer, consider the overly simplified two-layer (film) system analyzed by Liss and Slater [87] (figure 2-1). Within each film-like fluid layer, molecular diffusion processes are assumed to dominate the mass transport. Also, the bulk of each fluid is assumed well-mixed. For this scenario, only sparingly soluble gases, such as the environmentally relevant gases O₂, CO₂, and N₂, are considered. The flux across the air–water interface, F , is governed by Fick’s law for diffusion, which in one-dimensional form is

$$F = -D \frac{\partial C}{\partial z}, \quad (2.1)$$

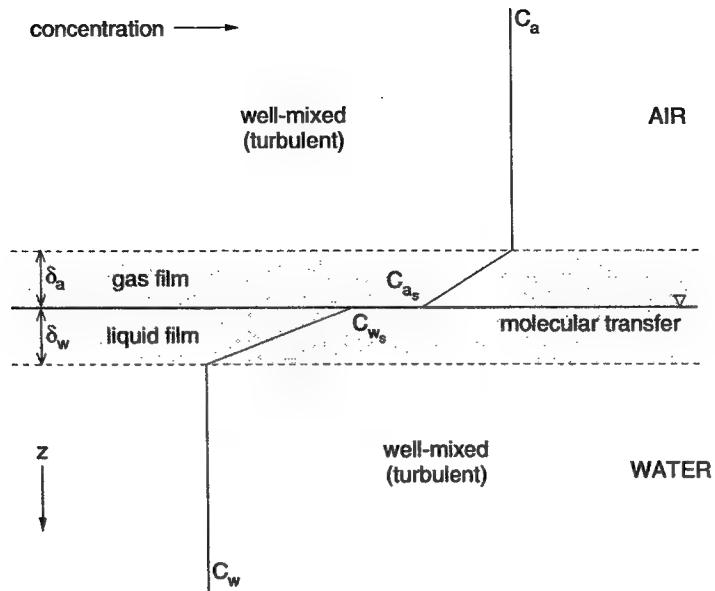


Figure 2-1: Sketch of air–water gas transport process. An idealized gas concentration profile is shown, where C_a is the bulk concentration in the air, C_w is the bulk concentration in the water, and C_{a_s} and C_{w_s} are the equilibrium airside and waterside concentrations at the air–water interface, respectively.

where C is concentration and z is positive downward. In this simple case, the flux can be written as

$$F = D_i \frac{\Delta C_i}{\delta_i}, \quad i = a, w, \quad (2.2)$$

where ΔC_i is the concentration difference across a layer, each layer having thickness δ_i . Since the flux also can be expressed as the product of a concentration difference and a mass transfer coefficient, or transfer velocity k , *i.e.*,

$$F = k_i \Delta C_i, \quad (2.3)$$

the gas-transfer velocity is often modeled as

$$k_i = D_i / \delta_i \quad (2.4)$$

for each layer. The D_i are the respective mass diffusivities of the gas in each fluid. At steady-state, equation 2.3 gives the mass flux balance through the interface,

$$F = k_a(C_a - C_{a_s}) = k_w(C_{w_s} - C_w). \quad (2.5)$$

Henry's law is used to relate the equilibrium air and water concentrations at the surface:

$$C_{a_s} = HC_{w_s}, \quad (2.6)$$

where H is the Henry's law constant ([equilibrium concentration in gas phase]/[equilibrium concentration of dissolved gas in liquid phase]). Eliminating C_{a_s} and C_{w_s} from equation 2.5 yields

$$F = \frac{C_a - HC_w}{1/k_a + H/k_w} = \frac{C_a/H - C_w}{1/k_w + 1/Hk_a}. \quad (2.7)$$

This can be rewritten as

$$F = K_a(C_a - HC_w) = K_w(C_a/H - C_w), \quad (2.8)$$

where

$$1/K_a = 1/k_a + H/k_w \quad (2.9)$$

and

$$1/K_w = 1/k_w + 1/Hk_a. \quad (2.10)$$

The quantities $1/K_i$ indicate the relative resistance to transport on either a gas-phase or liquid-phase basis. Writing equation 2.10 as a sum of two resistances,

$$R^w = r_w + r_a, \quad (2.11)$$

the relative contributions of the waterside and airside resistances to the total (liquid-based) resistance can be examined. For the waterside resistance,

$$r_w = \frac{1}{k_w} = \frac{\delta_w}{D_w}. \quad (2.12)$$

As a first approximation, the thickness of the waterside film layer will be assumed to scale as $D_w^{1/2}$ (see section 2.2.2). This gives

$$r_w \sim \frac{D_w^{1/2}}{D_w} \sim \frac{1}{D_w^{1/2}}. \quad (2.13)$$

Similarly, for the airside resistance,

$$r_a = \frac{1}{Hk_a} = \frac{\delta_a}{HD_a} \sim \frac{D_a^{1/2}}{HD_a} \sim \frac{1}{H} \frac{1}{D_a^{1/2}}. \quad (2.14)$$

Since, in general, $D_a \gg D_w$, and H is $\mathcal{O}(10)$ for slightly soluble gases, $r_w \gg r_a$. Thus, the transport is rate-limited by the waterside resistance, and the airside resistance can be considered negligible. This result is commonly accepted for all sparingly soluble gases that are not chemically reactive. This result equivalently implies that $k_w \ll k_a$, and that $\Delta C_w \gg \Delta C_a$, the latter condition implying that the total air–water concentration difference is dominated by the waterside gradient; with the concentration gradient in the air being slight, the simplifying assumption of $C_{a_s} = C_a$ is often made. In this case, the flux (equation 2.5) can be written in terms of *bulk* concentrations only:

$$F = k(C_a/H - C_w), \quad (2.15)$$

where $C_{a_s} = HC_{w_s}$ has been used. For highly soluble gases, such as methylamines, ammonia, and water vapor itself, the flux is limited by the airside resistance (McGillis *et al.* [94]).

Since the majority of the resistance to transfer across the air–water interface for the gases considered here lies in the waterside layer, the nature of this layer becomes crucial. The thickness of δ_w , which is inversely proportional to the rate of transfer, is determined by the diffusivity and the local hydrodynamic regime. The presence of turbulent mixing near the interface acts to thin this layer, steepening the concentration gradient, and enhancing transport. Complicating influences such as surfactants modify the near-surface flow and thereby affect the thickness of δ_w . In the next section, the thickness of the controlling waterside diffusion layer is explored in basic terms using boundary layer concepts.

2.2.2 Free-surface boundary layers

The transport of a passive, non-reacting scalar in a fluid medium is described by the convection-diffusion equation,

$$\frac{\partial C}{\partial t} + \mathbf{u} \nabla C = D \nabla^2 C, \quad (2.16)$$

where $C = C(\mathbf{x}, t)$ is the concentration of the scalar, \mathbf{u} is the three-component fluid velocity vector, and D reflects the diffusivity of the scalar in the fluid. For mass, C is simply the dissolved concentration of the species and D is the mass diffusion coefficient of the species in a particular solute. Similarly for heat (*viz.*, temperature, T), C is equivalent to $\rho c_p T$, where ρ is the density of the fluid and c_p is the specific heat of the fluid at constant pressure; the diffusivity becomes $D = \kappa / \rho c_p$, where κ is the thermal conductivity in the fluid. Considering the reduced problem of a planar air-water interface in only two dimensions (x , horizontal and z , vertical), the concentration of a non-reactive gas in water is governed by the simplified convection-diffusion equation,

$$\frac{\partial C}{\partial t} + u \frac{\partial C}{\partial x} + w \frac{\partial C}{\partial z} = D \left(\frac{\partial^2 C}{\partial x^2} + \frac{\partial^2 C}{\partial z^2} \right), \quad (2.17)$$

where D is now the diffusivity of the gas in water. In this equation, the first term of the left-hand-side is just the Eulerian time rate of change of the gas concentration. The second and third terms describe the convective transport of the gas, and the right-hand-side is the Fickian diffusion term. The role of convection in the transport is determined by the solutions of the Navier-Stokes equations for the components of the water velocity:

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + w \frac{\partial u}{\partial z} = -\frac{1}{\rho} \frac{\partial p}{\partial x} + \nu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial z^2} \right) \quad (2.18)$$

$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + w \frac{\partial w}{\partial z} = g - \frac{1}{\rho} \frac{\partial p}{\partial z} + \nu \left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial z^2} \right) \quad (2.19)$$

where g is the acceleration due to gravity, p is the fluid pressure, and ν is the kinematic viscosity (μ/ρ). Thus, the rigorous specification of $C(\mathbf{x}, t)$ involves the solution of the nonlinear Navier-Stokes equations to determine the flow field that then can be used to solve

the convection-diffusion equation for the concentration. In the absence of such a solution, a dimensional scaling approach to the governing equations is useful for the present discussion. To simplify the analysis, the case of a steady, two-dimensional stagnation flow at a plane free surface is studied. Stagnation flows have been used in a number of other studies as a model for a surface renewal process important in gas exchange (*e.g.*, Brumley and Jirka [14]; Chan and Scriven [17]; Csanady [22]). The use of model stagnation flows is predicated on the assumptions that the flow adjacent to the interface is purely irrotational (viscous effects negligible) with a linear vertical velocity profile and that the diffusive transport is solely vertical. For this model, the flow field can be represented by the stream function

$$\psi = axz, \quad (2.20)$$

where a is the divergence strength ($a > 0$) having units of [velocity/length] and $z = 0$ is the free surface with z positive downward. The velocity components follow as

$$u = ax, \quad w = -az. \quad (2.21)$$

The equation for the gas concentration (equation 2.17) can be easily scaled by considering only the situation along the stagnation streamline $x = 0$, where there exists symmetry about the z -axis, $\partial C / \partial x = 0$. A further assumption that $\partial^2 C / \partial x^2 \ll \partial^2 C / \partial z^2$ also can be made based on the arguments of Csanady [22]. This leaves

$$w \frac{\partial C}{\partial z} = D \frac{\partial^2 C}{\partial z^2}. \quad (2.22)$$

Nondimensionalizing the above equation using $\tilde{z} = z/\delta_s$, $\tilde{w} = w/(a\delta_s)$ and $\tilde{C} = C/C_*$, where δ_s loosely represents a scalar diffusion layer thickness, yields

$$\tilde{w} \frac{\partial \tilde{C}}{\partial \tilde{z}} = \left(\frac{D}{a\delta_s^2} \right) \frac{\partial^2 \tilde{C}}{\partial \tilde{z}^2}. \quad (2.23)$$

If the scales have been chosen correctly, the quantity in parentheses is $\mathcal{O}(1)$, and $\delta_s^2 \sim D/a$.

Introducing a characteristic length scale L and the Schmidt number,

$$Sc = \frac{\nu}{D}, \quad (2.24)$$

provides the result

$$\frac{\delta_s}{L} \sim Re_a^{-1/2} Sc^{-1/2}, \quad (2.25)$$

where $Re_a = aL^2/\nu$ is a Reynolds number based on the divergence. In very general terms, $\delta_s \sim \mathcal{O}(100 \mu\text{m})$, or about $1/10^{\text{th}}$ the scale for momentum, which would typically scale as $\delta_m \sim L Re_a^{-1/2}$. Using the expression for the waterside transfer velocity given in equation 2.4, it is found,

$$k \sim (aD)^{1/2}, \quad (2.26)$$

or, in terms of the Reynolds and Schmidt numbers,

$$\frac{k}{aL} \sim Re_a^{-1/2} Sc^{-1/2}, \quad (2.27)$$

where aL has been used as a characteristic velocity scale. The $Sc^{-1/2}$ dependence for the transfer velocity is commonly accepted for *clean* air–water interfaces, where the free surface allows a free-slip boundary condition. If the interface is taken to be a *rigid* wall, a viscous boundary layer forms, giving rise to a “Hiemenz flow” (Schlichting [117]). Csanady worked out the form of the gas-transfer velocity for this situation:

$$\frac{k}{aL} \sim Re_a^{-1/2} Sc^{-2/3}. \quad (2.28)$$

The only difference between equations 2.27 and 2.28 is the magnitude of the exponent of the Schmidt number (hereafter denoted n). That exponent therefore will depend upon the nature of the interface—free-slip ($n = 1/2$) or no-slip ($n = 2/3$). Actual air–water interfaces lie somewhere between the two bounds due to surfactant effects (discussed in the next section). The importance of the Schmidt number exponent n , in addition to its role in equations 2.27 and 2.28, is its use to scale transfer velocity results from one gas to infer

the transfer velocity for another. This is done using the expression

$$\frac{k_1}{k_2} = \left(\frac{Sc_1}{Sc_2} \right)^{-n}. \quad (2.29)$$

Equation 2.29 is of critical importance to the recent efforts of using heat as a proxy tracer for gas. As seen from equation 2.16, the transport of heat and the transport of mass are analyzed in the same manner. Therefore, it is possible to imagine heat as a gas with a very high diffusivity, or low Schmidt (Prandtl) number. With a measurement of the transfer velocity for heat, equation 2.29 can then be used to infer the transfer velocities for other gases of interest. However, because the Prandtl number is roughly two orders of magnitude smaller than the Schmidt number for gases such as CO₂ and O₂, the scaled transfer velocity is very sensitive to the Schmidt number exponent n . Another point that arises when considering a measurement of interfacial heat transfer involves the airside and waterside temperature gradients. Temperature behaves differently than the slightly soluble gases considered here. Because the heat capacity of water is more than three orders of magnitude greater than that for air, it is found that the dominant contribution to the total air–water temperature difference occurs in the air. Therefore, on a waterside basis, the flux of heat is dictated by difference between the bulk water temperature and the surface temperature. Accurate and precise measurements of water surface temperature are challenging.

At present, innovative techniques pioneered by Haussecker *et al.* [50] are exploring the temporal decay of heat within small surface fluid patches at the air–water interface. Using infrared imagers, these methods track the heat dissipation through the surface layer, in analogy to the diffusion of gas through this layer. Through measurement of the time constant of the decay (< 100 seconds) and invoking conceptual models, the transfer velocity for heat can be estimated. Initial measurements with these methods have yielded results that are in agreement with other more traditional techniques.

2.3 Surface film effects

2.3.1 Surface tension

Of the many properties of an interface, surface tension is perhaps the most widely recognized. The concept of surface tension is actually a construct of convenience that is used in lieu of the surface free energy (Adam [1]). That there exists free energy associated with a surface is revealed by its natural tendency to contract spontaneously. This is the reason for the spherical shape assumed by bubbles and droplets, the beading of droplets on a solid surface, and the phenomenon of capillarity. Surface free energy, or surface tension, can be rationalized by considering the molecular attraction forces in a liquid. In the fluid bulk, individual molecules are surrounded by neighboring molecules on all sides, and therefore experience attractive forces omnidirectionally, which, on average, are uniform in all directions. However, at a surface, molecules tend to be pulled inwards (toward the bulk) and laterally because the forces of attraction outwards are much less owing to the fewer number of molecules outward (*i.e.*, the molecular bonds formed between surface molecules and their surface/subsurface neighbors are enhanced). Consequently, the area of the surface will diminish and the fluid will contract until the smallest possible surface area for the given fluid volume is reached (hence the spherical shape in the absence of external forces). It therefore requires a finite amount of work to bring molecules from the bulk to the surface against these inward attractive forces. The required work is proportional to the increase in surface area. This work is expressed in terms of energy per unit area, or force per unit length and is called surface tension, denoted σ . An interfacial tension will exist whenever two immiscible fluids are in contact. The magnitude of the interfacial tension will depend on the physical properties of the two fluids, and in general is a function of pressure, temperature, and composition (Sabersky *et al.* [114]).

There are several different methods for experimentally determining surface tension (*e.g.*, the tensiometry methods recently reviewed in Lapham *et al.* [80]). The technique used throughout this thesis is known as the Wilhelmy plate method (Davies and Rideal [25]). In this technique, a thin plate is suspended vertically from a weighing balance and is gently brought in contact with a free surface (refer to figure 2-2). To insure proper wetting, the

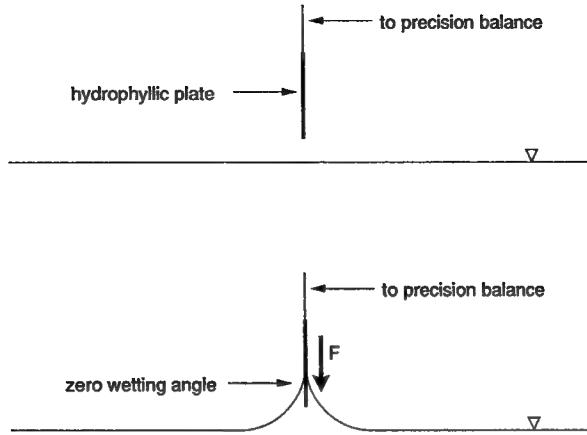


Figure 2-2: Diagram of the Wilhelmy plate technique. In the upper sketch, the plate is shown suspended from a precision balance. In the lower sketch, the plate is in contact with the fluid and a zero wetting angle has formed. The force F measured by the balance can be related to the surface tension.

plate can be finely roughened, although Lapham *et al.* point out that such roughening is a current topic of controversy. When the fluid wets the plate, an additional force will be detected by the balance, this force being related directly to the surface tension by the commonly used expression,

$$F = \sigma P \approx 2\sigma L, \quad (2.30)$$

where F is the force and P is the plate perimeter, which for very thin plates is well approximated by twice the plate length L . Higher order corrections can be made for plate end effects and buoyancy effects. Equation 2.30 is only valid when the contact angle between the fluid and the plate is zero (perfect wetting). In all other cases, a multiplicative cosine factor must be incorporated to account for a finite wetting angle. The units for surface tension used in this thesis are millinewtons per meter (mN/m), which are equivalent to dynes/cm, the commonly used unit in the CGS system.

Previously obtained experimental data using the Wilhelmy plate technique were provided to the present author that furnished the interfacial air–water tension of ultra-high pure (UHP) water as a function of temperature (Nelson [103]). Here, UHP water is used to describe type I ultra-pure water, re-distilled out of an alkaline potassium permanganate

solution. A linear, least-squares regression of the data yielded the following expression for the reference surface tension of UHP (“clean”) water:

$$\sigma_o(T) = 75.629 \frac{\text{mN}}{\text{m}} - 0.137 \frac{\text{mN}}{\text{m}^\circ\text{C}} T, \quad (2.31)$$

with T in degrees Celsius.

As an alternative to equation 2.30, an empirical relationship can be operationally obtained to relate the measured force to the surface tension of the liquid (*i.e.*, $\sigma = \alpha F$, where α is a constant for a given balance-plate setup). Such an approach was taken here. For all surface tension work in this thesis, a finely roughened platinum plate of 0.13 mm thickness was used. Prior to each measurement, the plate was cleaned by rinsing with a 50:50 methanol/water solution and flamed in a propane/air flame. To determine the constant α , 23 measurements of F were made for UHP water. The temperature of the UHP water was measured and used to determine the reference surface tension using equation 2.31. These data were analyzed to estimate α . Based on the 95% confidence intervals, measurement uncertainty of surface tension using this approach was estimated to be $\pm 0.28\%$.

2.3.2 Surfactants

It is well known that the surface tension of an interface is notably reduced by the presence of a surface-active agent, or surfactant. Typically large molecules, surfactants are bipolar in structure, being comprised of a hydrophobic (water-“fearing”) part and a hydrophilic (water-“loving”) part. Consequently, surfactants seek an equilibrium state at an interface between aqueous and non-aqueous phases—*e.g.*, the air–water interface. Figure 2-3 shows a surfactant molecule of oleic acid ($\text{C}_{18}\text{H}_{34}\text{O}_2$). Oleic acid molecules are composed of a hydrophobic hydrocarbon chain, or ‘tail,’ and a hydrophilic polar $-\text{COOH}$ acid group, the ‘head.’ The low degree of polarity in these molecules renders them insoluble and they adopt a preferential orientation on the surface (taken here to be water) as shown in the lower sketch of figure 2-3. Highly insoluble surfactants form insoluble monolayers at the surface. Molecules with shorter hydrocarbon chains and/or head groups with greater polarity may be soluble in the aqueous phase.

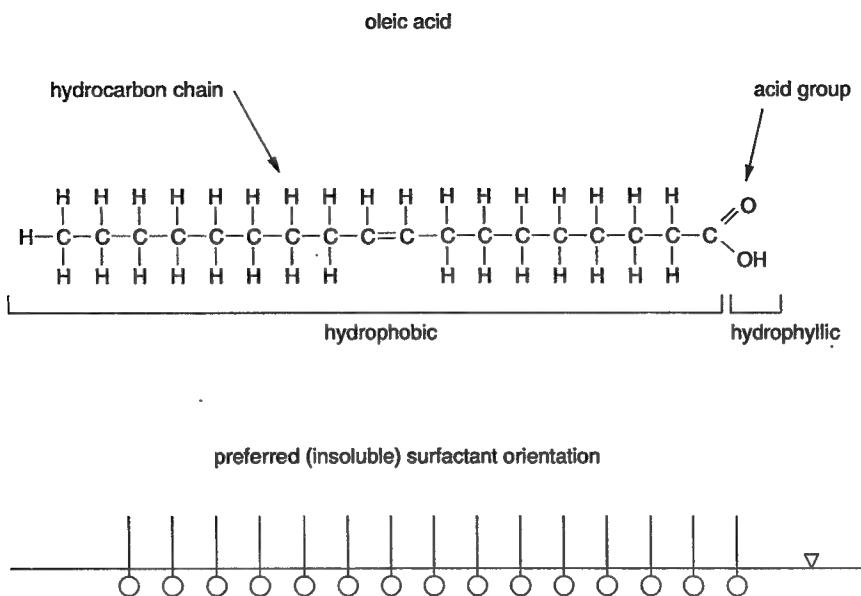


Figure 2-3: Molecule of oleic acid ($C_{18}H_{34}O_2$) and preferred insoluble surfactant orientation. In the lower sketch, the \circ symbol represents the polar 'head,' and the — symbol corresponds to the hydrocarbon 'tail.'

When a surface-active film is present at a free surface, it will tend to lower the surface free energy, thus lowering the surface tension. The amount by which the surface tension is reduced is termed the surface pressure π , and is defined as

$$\pi = \sigma_o - \sigma, \quad (2.32)$$

where σ_o is the surface tension under surfactant-free equilibrium conditions. The dependence of the surface tension on surfactant concentration for soluble and insoluble surfactants may be determined experimentally using a setup such as a Langmuir trough (Davies and Rideal [25]). Data from such experiments are typically reported as adsorption isotherms, or π -A isotherms, which plot quasi-static surface pressure versus film area or area per molecule. Relevant surface pressure isotherms for this thesis work are presented in section 7.3.2.

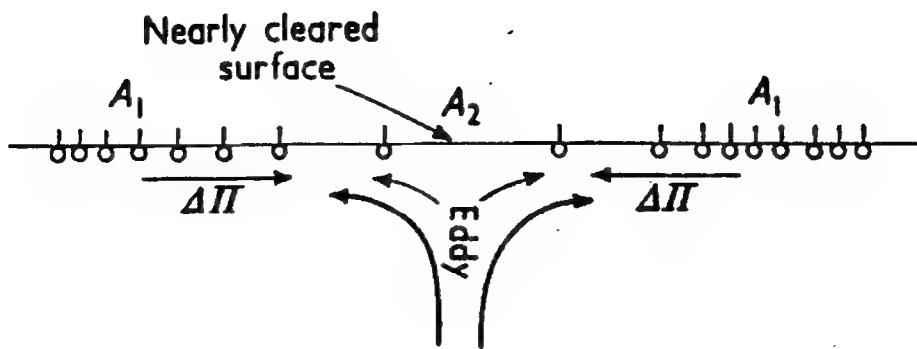


Figure 2-4: Conceptualization of surface renewal by a turbulent eddy in the presence of a surface film. The eddy upwells “clean” fluid to the surface, while the differential spreading pressure, $\Delta\pi$, opposes the transport of new fluid to the interface. Source: Davies and Rideal [25].

2.3.3 Marangoni forces and surface elasticity

Based on π -A isotherm data, the surface tension of a surfactant-adsorbed interface will depend on the surfactant concentration. Spatial variations in concentration therefore result in variations in surface tension that manifest themselves as surface tension gradients. Free-surface temperature variations will produce a similar effect. Gradients in σ , or equivalently, in π , can drive surface flows as well as have profound implications for the near-surface bulk flow below. These effects were observed by Marangoni (1871), and are generally referred to as Marangoni effects or forces. The Marangoni effect can be thought of in several different ways. The sketch in figure 2-4 illustrates one such way. While a turbulent eddy impinges at the water surface, it redistributes surfactant on the surface. Since the concentration of surfactant is greater away from the point of upwelling (A_1), the surface tension is correspondingly less in this region. At the point of upwelling (A_2), the surface has been cleared of nearly all surfactant and the surface tension is high. A gradient in surface tension therefore is created that drives a flow back inward, opposing the eddy motion. In this manner, surfactants and the Marangoni effect damp surface renewal by resisting surface divergent motions.

A more formal approach to understanding Marangoni forces is to consider how these forces modify the free-surface tangential boundary condition. This boundary condition

is crucial for near-surface diverging flows and capillary wave dynamics. At a truly free, smooth surface, void of surface tension gradients and ignoring airside effects ($\mu_w \gg \mu_a$), the simplified stress boundary condition tangential to the interface states

$$\mu \frac{\partial u}{\partial z} = 0 \quad (2.33)$$

for unidirectional flow in two dimensions. However, when a surface film is present and the flow has created a situation where $\partial\sigma/\partial x \neq 0$, the boundary condition becomes

$$\mu \frac{\partial u}{\partial z} = \frac{\partial \sigma}{\partial x}. \quad (2.34)$$

The surface tension gradient therefore is balanced by the viscous stress of the fluid. The presence of a non-zero viscous stress at the interface generates a highly dissipative viscous boundary layer in which flow damping takes place. Such flow damping acts to thicken the aqueous mass boundary layer, and can retard gas exchange significantly. The influence of these additional forces can become quite complex when the stress balance at the surface modifies the hydrodynamic flow in the bulk, which in turn redistributes the surface film and leads to new surface stresses. In effect, there is a closed-loop interaction between the surface forces and the near-surface bulk hydrodynamics (Tsai and Yue [133]).

Under static or quasi-static conditions, a surfactant-adsorbed surface will exhibit a surface elasticity often called the equilibrium, or Gibbs, elasticity. This elasticity describes the resistance of the surface to in-plane compression/dilation of the surface. The static elastic modulus, ϵ_o , is defined as

$$\epsilon_o = \frac{d\sigma}{d \ln A}, \quad (2.35)$$

which relates the dilational strain (A is the area of a surface element) to the surface stress. This quantity is helpful in understanding the compressional behavior of a surface when the rates of compression/dilation are on the order of an hour or greater (Bock and Frew [10]). Data from adsorption isotherms are typically used to estimate ϵ_o .

Under certain dynamical conditions, a surface may exhibit rate-dependent resistance to in-plane compression and dilation. The passage of a wave, for instance, will compress a

surface film near the front of the wave, while dilating it on the downward moving surface behind the wave. When the time scales associated with the dynamic perturbations are comparable to the time scales involved with the film adsorption/desorption and molecular re-orientation processes, the elastic modulus is modeled as having a viscous component. Additionally, if the rate of surface strain is non-zero, the elasticity of the film may deviate from its equilibrium value. The result is a dynamic viscoelastic modulus (typically represented in complex form) for compression/dilation that is composed of a surface dilational elasticity and a surface dilational viscosity. When the dilational modulus is real-valued the surface acts as a purely elastic film. For pure water, the viscoelastic modulus vanishes entirely and the interface is stress-free. However, the dilational moduli for typical samples of sea water are found to be complex owing to the presence of a rich mixture of natural surfactants.

A similar viscoelastic modulus exists for surface shear that governs the response of the surface to dynamic deformations that alter the shape of elements of free-surface area in a sense that is not purely dilational. The viscoelastic dilational modulus and the viscoelastic shear modulus comprise the formal surface stress tensor that is found in classical rheological treatments (Goodrich [44]). These four rheological coefficients govern the response of a surface undergoing changes in area (dilation) and/or shape (shear). Considering only dilational/compressional effects, the viscoelastic modulus can be written as

$$\epsilon = \frac{\partial \sigma}{\partial \ln A}, \quad (2.36)$$

which will be a function of the time scales of the surface motion. Rearranging this expression as in Lucassen-Reynders and Lucassen [90], the surface tension gradient can be rewritten as

$$\frac{\partial \sigma}{\partial x} = \epsilon \frac{\partial^2 \xi}{\partial x^2}, \quad (2.37)$$

where ξ is the horizontal deviation of the surface from equilibrium. Equation 2.34 then becomes

$$\mu \frac{\partial u}{\partial z} = \epsilon \frac{\partial^2 \xi}{\partial x^2}. \quad (2.38)$$

Expressing the free-surface boundary condition in this manner clearly reveals how the sur-

face strain and the viscoelastic modulus enter the stress balance at the interface.

Measurement of the static elasticity is possible using a Langmuir trough and equation 2.35. Measurements of the dynamic viscoelastic modulus, on the other hand, are not as straightforward, requiring more sophisticated methods and dedicated experimental setups. Therefore, a measurement of the dynamic elasticity in the field, or even in a laboratory tank used for unrelated hydrodynamic studies, is exceptionally difficult. As such, the range of viscoelastic properties for natural water surfaces is not well known, and the measurement of the elastic modulus in the field is currently an active area of research. In chapter 7, static elasticities for representative samples of the water surfaces studied in that chapter are presented. However, the relevance of these measurements to the observed hydrodynamics and/or gas transfer is dubious since the measurements were not made in the laboratory tank where the experiments took place. The data are only included to illustrate important differences between the samples. In lieu of elasticity, *in situ* measurements of the surface pressure are used to qualify surface cleanliness throughout this thesis. These measurements were made in the laboratory tank and were obtained in conjunction with the hydrodynamic and gas-transfer measurements. This is a step forward in terms of proper consideration of the role surface films play in interfacial studies.

2.4 Models of air–water gas transport

Modeling air–water gas transfer is particularly important to environmental scientists today. Developing models of transport across the air–sea interface based on experimental data is necessary for climate modeling and heat and mass budgets. The present challenge is not only to develop models that properly capture the relevant physics, but also to cast these models such that the model inputs are actually measurable quantities that have physical meaning. Most gas-transfer models can be generalized as being either conceptual, eddy diffusivity, or hydrodynamic models (Brumley and Jirka [14]). In the next few sections, a representative selection of examples of these various model types are discussed. In chapter 7, some of these models are explored.

2.4.1 Conceptual models

Conceptual models typically fall into the class of either thin-film models or surface renewal models, with many variant forms as well. The thin-film model (Whitman [143]) likens the situation to smooth flow near a wall where there exists a thin stagnant layer at the interface and the fluid below is well mixed. The physical picture is identical to the scenario depicted in figure 2-1. The (waterside) gas-transfer velocity is given simply in terms of the gas diffusivity and the thickness of the waterside diffusion layer δ ,

$$k = \frac{D}{\delta}. \quad (2.39)$$

Studying a clean air–water interface, Davies and Rideal [25] have shown through a more detailed analysis that $k = D/2\delta$ where δ is still taken to be the molecular diffusion layer thickness and is considered stagnant for mathematical purposes. Through turbulent boundary layer scaling, they found

$$\delta = \left(\frac{D\sigma_e}{0.4\rho u_o^3} \right)^{1/2}, \quad (2.40)$$

where u_o is a velocity scale for the near-surface bulk turbulence and σ_e is an equivalent surface tension that includes gravitational restoring effects. The resulting expression for the transfer velocity is

$$k = 0.32D^{1/2}u_o^{3/2}\rho^{1/2}\sigma_e^{-1/2}. \quad (2.41)$$

While greatly simplified, the film model incorporates an essential feature of the true process of interfacial transport: gas must first diffuse away from the interface via molecular processes before turbulent convective effects can become important.

Higbie [54] and Dankwerts [23] both proposed transport models based on the ideas of surface renewal and penetration. In these formulations, the free surface is assumed to be populated with an array of surface parcels that are periodically replenished with bulk fluid elements by the turbulent flow below. Transport into an individual surface parcel occurs during the time it resides at the surface. Assuming Fickian diffusion for a parcel, the transfer velocity was found to scale as $k \sim (D/\tau)^{1/2}$, where τ represents a characteristic surface renewal time, or “eddy age.” Higbie postulated that the eddy age τ was identical

for all surface elements. This yielded the classical result,

$$k = 2 \left(\frac{D}{\pi \tau} \right)^{1/2}. \quad (2.42)$$

Dankwerts did not assume a constant renewal time, but rather proposed a probability distribution of renewal times. In this manner, Dankwerts obtained

$$k = \left(\frac{D}{\tau} \right)^{1/2}. \quad (2.43)$$

The expressions of Higbie and Dankwerts are virtually identical. Furthermore, both require an estimate of the surface renewal time to predict the transfer velocity. The relevance of this time scale to a measurable quantity is not necessarily evident.

2.4.2 Eddy diffusivity models

In the eddy diffusivity approach, which is a standard in turbulence studies, an effective diffusivity due to the flow itself is introduced, and the flux is given as

$$F = -(D + D_e) \frac{\partial C}{\partial z}. \quad (2.44)$$

Applying this concept to mass transfer, previous works have typically assumed an eddy diffusivity $D_e \sim z^m$, $m > 0$, yielding multiscale models. The result is $k \sim D^n$, where $0 < n < 1$ depending on the value for m , which is a function of the interfacial condition. For a truly free surface, $D_e \sim z^2$, and

$$k \sim Sc^{-1/2}. \quad (2.45)$$

For a smooth solid surface, $D_e \sim z^3$, and

$$k \sim Sc^{-2/3}. \quad (2.46)$$

These last two expressions mimic the results given in equations 2.27 and 2.28. In a simple eddy diffusivity model, Ledwell [82] used scaling arguments to deduce a limiting result for

a clean interface,

$$k \rightarrow \frac{2}{\pi} (Da)^{1/2}, \quad (2.47)$$

where a is a scale of the turbulent surface divergences (Ledwell suggested a root-mean-square value). This result points to the potential significance of surface divergence to interfacial gas exchange. Turbulent eddies near the free surface are important to air–water gas exchange because these eddies bring fresh fluid near to the interface for diffusion. Frequent vertical transport of fresh fluid toward the surface results in a reduction of the mass boundary layer thickness. In this way, the vertical velocity fluctuations near the interface are considered vital to gas-transfer control. The importance of the surface divergence can be readily seen if the assumption is made that the vertical velocity fluctuations near the surface, $w'(z)$, which bring fluid toward the interface for diffusion, are responsible for gas-transfer. Treating water as incompressible, by continuity,

$$\frac{\partial w'}{\partial z} = -\nabla_h \cdot \mathbf{v}', \quad (2.48)$$

where $\nabla_h \cdot \mathbf{v}'$ is the two-dimensional velocity fluctuation divergence in a plane locally tangent to the fluid interface. A Taylor series expansion yields an expression for the vertical velocity fluctuations:

$$w'(z) = w'(0) - (\nabla_h \cdot \mathbf{v}')_0 z - \left[\frac{\partial}{\partial z} (\nabla_h \cdot \mathbf{v}') \right]_0 \frac{z^2}{2} + \dots, \quad (2.49)$$

where the subscript ‘0’ indicates evaluation at $z = 0$. In a reference frame moving with the interface, $w'(0)$ vanishes at the free surface. Then, for small z , higher order terms can be neglected and this reduces to

$$w'(z) \approx - \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right)_0 z, \quad (2.50)$$

where z is much less than the characteristic lengthscale of the horizontal velocity field near the boundary (Ledwell [82]). Brumley and Jirka [14] have argued that equation 2.50 is a valid approximation within one Kolmogorov scale of the surface, where the suppression of horizontal vorticity results in irrotational, stagnation-like flow. Expressing $w'(z)$ in such manner illustrates the potential direct link between gas transfer and the surface divergence. The velocity fluctuations responsible for thinning the mass boundary layer and increasing

the interfacial gas flux are closely related to the divergence of the velocity field at the surface. A more sophisticated eddy diffusivity model developed by Coantic [19] included the effects of wind-driven shear and capillary waves. The result was an expression for the gas-transfer velocity similar to that of Ledwell, but instead of a surface divergence scale, the model predicted a transfer velocity that is proportional to the correlation between the divergence and the near-surface shear.

2.4.3 Hydrodynamic models

Hydrodynamic models search for solutions to the convection-diffusion equation using idealized flow motions. Assuming a flow field composed of a regular array of roll cells at the surface, a *large eddy* model was developed by Fortescue and Pearson [33]. Taking the approach that the relatively large scale eddies are most important to interfacial gas exchange, the predicted transfer velocity was determined to be

$$k = 1.46 \left(\frac{DQ}{L} \right)^{1/2}, \quad (2.51)$$

where Q is a root-mean-square (RMS) bulk turbulent velocity scale and L is the integral lengthscale of the underlying bulk turbulent flow. A similar roll cell model, which considered the small scale eddies as driving gas transport, was proposed by Lamont and Scott [78]. This *small eddy* model, together with an estimate of the turbulent energy spectrum, yielded the result

$$k = 0.4D^{1/2} \left(\frac{\epsilon}{\nu} \right)^{1/4}, \quad (2.52)$$

where ϵ is the turbulent dissipation rate. This result is of the same form as found by Kitaigorodskii [72] who studied the problem of gas transport due to turbulent patches generated by breaking waves.

Another group of hydrodynamic models has evolved from the use of a stagnation flow to specify the flow field. Chan and Scriven [17] solved the convection-diffusion problem assuming a steady stagnation flow and found the following expression for the instantaneous transfer velocity:

$$k(t) = \sqrt{\frac{2}{\pi}} \frac{(Da)^{1/2}}{\sqrt{1 - e^{-2at}}}, \quad (2.53)$$

where a is the surface divergence and t is the exposure time. This expression is very similar to that of Ledwell, with the additional exponential factor to account for the time dependent behavior of diffusion at the interface. Using the same stagnation flow and ignoring any time-dependence, Csanady [22] found $k = \sqrt{2/\pi} (Da)^{1/2}$, which is identical to the Chan and Scriven result for $at \gg 1$. Brumley and Jirka [14], using certain concepts from the Chan and Scriven approach, have proposed a model that is based on surface divergence *fluctuations*. Others have also considered the role of surface divergence fluctuations, or the surface divergence spectra. McCready *et al.* [92] investigated the role of the spectra of the vertical velocity fluctuations through numerical simulations, Back and McCready [7] incorporated short wavelength wave effects, and Gulliver and Tamburrino [46] have attempted to explore these results with laboratory experiments. These works have focused on an expression of the gas-transfer velocity that uses spectral properties of the divergence (McCready *et al.* [92]):

$$k^+ = c Sc^{-1/2} \left(\overline{a^2} \right)^{1/4}, \quad (2.54)$$

where k^+ is a dimensionless transfer velocity, scaled using the friction velocity, c is a multiplicative constant, and

$$\overline{a^2} = \int_0^\infty S_a(\omega) d\omega, \quad (2.55)$$

with $S_a(\omega)$ representing the spectral density function for the surface divergence. The divergence and its spectra appearing in the previous two expressions are made dimensionless using the friction velocity and the kinematic viscosity. Since the square root of equation 2.55 is the RMS of the divergence, the expression for the k^+ in equation 2.54 indicates that the transfer velocity is proportional to the RMS of the surface divergence to the 1/2-power. This agrees with the previously discussed models incorporating the surface divergence.

2.4.4 Other parameterizations

The models outlined above, while potentially helpful in revealing important aspects of the air–water transfer problem, are somewhat lacking in applicability in the field (*i.e.*, over the ocean). The choices of appropriate velocity scales and lengthscales is unclear, surface divergence has yet to be measured in the field, and surfactant influences are always a com-

plicating effect. Consequently, simple empirical relationships that attempt to predict the gas-transfer velocity using measurable quantities are the current state-of-the-art in field studies. Liss and Merlivat [86] have proposed a three-regime piecewise *linear* parameterization for the transfer velocity based on the 10 meter wind speed (U_{10}). Wanninkhof [136] has proposed a *quadratic* relationship between the transfer velocity and U_{10} , and Wanninkhof and McGillis [137] have proposed a *cubic* form. Each parameterization includes the $Sc^{-1/2}$ factor. Jähne *et al.* [64] have advanced a popular relationship that depends on the friction velocity:

$$k = \frac{1}{\beta} Sc^{-n} u_*, \quad (2.56)$$

where β is a constant dependent upon the hydrodynamic regime and u_* is the waterside friction velocity. In the many forms of this parameterization, the exponent n has fallen between 1/2 and 2/3, and β has ranged from 11 to 16. In more recent developments, experiments in wind-wave tanks have shown an improved relationship between the transfer velocity and surface roughness. Specifically, it has been found that the mean square slope correlates with the transfer velocity along a unique curve that is independent of the surface cleanliness (Frew [34]). The good correlation between the two is likely rooted in the mutual sensitivity of wave amplitude (damping) and surface renewal processes to surface film effects. Predictions based on wind speed or friction velocity are unable to capture such effects.

2.4.5 A model of gas transport including surface film effects

None of the models thus far described include any direct effect of the presence of a surfactant. The above models were developed for the limiting case of a perfectly clean air–water interface. To date, there is a paucity of modeling efforts that incorporate the important role of surface films on gas transfer. In two early works, Davies [24] and Davies and Rideal [25] studied the effect of a surface film on damping turbulent eddies at a free surface. In the analysis, the stress balance at the air–water interface resulting from a single eddy impinging at the surface was modeled using scaling arguments. In the absence of a surface film, the dynamic “thrust” of the eddy is balanced by the restoring force due to surface tension

(gravitational effects are ignored here):

$$\rho u_o^2 = \frac{2\sigma}{\mathcal{R}}, \quad (2.57)$$

where u_o is the characteristic velocity of the bulk turbulent eddies and \mathcal{R} is the minimum radius of curvature of the deformed interface due to the eddy. However, for a film-covered surface with a surface pressure π , a new term due to the resistance to surface clearing is introduced,

$$\rho u_o^2 = \frac{2\sigma}{\mathcal{R}} + \frac{\pi}{\delta_m}, \quad (2.58)$$

where δ_m is the viscous sublayer. By relating \mathcal{R} to δ_m , solutions for the viscous sublayer thickness can be found for clean and film-covered surfaces. Furthermore, Davies used the Schmidt number to relate the diffusive sublayer thickness, δ , to δ_m and then employed $k = D/\delta$ to calculate a transfer velocity. The effect of non-zero surface pressure π is to increase the thickness of the viscous sublayer, and in turn, the diffusive sublayer, thereby reducing the transfer velocity. Davies calculated an expected reduction of the gas-transfer velocity for a surfactant surface relative to that for a clean surface in terms of the surface pressure π :

$$\frac{k_{\text{film}}}{k_{\text{clean}}} = \left\{ \left[1 + \frac{0.25}{u_o \sigma \nu \rho} \left(\frac{\pi \mu b u_o}{\pi + \mu b u_o} \right)^2 \right]^{1/2} + \frac{0.5}{(u_o \sigma \nu \rho)^{1/2}} \left(\frac{\pi \mu b u_o}{\pi + \mu b u_o} \right) \right\}^{-1}, \quad (2.59)$$

where b is a constant. For the above result, k_{clean} is given by an expression that is effectively equivalent to equation 2.41. In addition, in deriving this result, Davies included viscous damping near the partially immobilized interface through the term $(\mu b u_o)$. In the limit of large π , the effect on the transfer velocity becomes dominated by viscous effects; for very small π , viscous effects become less critical. The Davies model is explored briefly in chapter 7 using measurements of turbulence-driven gas-transfer in the presence of surface films.

2.4.6 Summary discussion

Air–water mass transport is a process that is largely dependent upon the hydrodynamics very near the air–water interface. There is little question that near-surface turbulence and surface renewal events are key to enhancing the transport. Therefore, models or parameterizations that attempt to capture the effects of these mechanisms are likely to be successful—when surface film effects are not present. Additionally, a parameterization such as one involving wave slope, while showing good success, has the potential for drawbacks: a relationship between gas transfer and wave slope only implicitly incorporates surfactant effects and lacks physical interpretation. The slope of the waves is not the responsible mechanism for gas transfer, but rather is a proxy for another process(es) that is controlling. The fact that wave slope and gas transfer show a favorable relationship for a range of surface conditions is due to the similar effects surfactants have on both transport and wave slope. The surfactant-induced surface stress that results in a thicker mass boundary layer is also responsible for wave damping.

An association between surface divergence and mass transfer is expected to yield a more meaningful relationship, because as has been discussed, surface divergence can be physically linked to the gas-transfer process. In addition, as depicted in figure 2-4, surface films will impact surface divergences, and based on a simple analysis, only small amounts of surfactant (surface pressure) are needed to effect large changes in the surface dynamics. Therefore, there should exist a connection between surface divergence, gas transfer, and surfactants. Connections between these three elements are explored in chapter 7.

Chapter 3

Performance of digital image velocimetry processing techniques

3.1 Introduction

In this chapter, the performance of a number of simple digital particle image and particle tracking velocimetry techniques (DPIV and PTV) is examined. These techniques are used throughout this thesis, forming the backbone of almost all of the experimental studies presented. In essence, these techniques are quantitative flow visualization methods that process imagery of a particle-seeded flow to extract velocity information. While the digital versions of these flow measurement techniques have been around since the early 1990s, the roots of these methods date back several decades. Explanations and discussions of these techniques can be found in Adrian [2][3], Agüí and Jimenez [4], McKenna [96], Westerweel [139], and Willert and Gharib [144], to name only a few. In this work, the focus has been restricted to the more common approaches to DPIV and PTV processing found in practice today. Other more sophisticated techniques such as those of Huang *et al.* [59], Okamoto *et al.* [104], Kumar and Banerjee [76], Quénod *et al.* [107], and Song *et al.* [124] are not considered, but are noted for their significance in advancing the state-of-the-art in DPIV and PTV processing techniques. One reason for this focus on more standard processing methods is that many commercial image velocimetry packages available today are based upon traditional forms of analysis (Fast Fourier Transform techniques, for example). Understanding the limitations

and expected performance of these types of processing codes is therefore necessary. A second reason for this focus is the relative simplicity of these methods and the lack of intensive algorithm development required for implementation.

In addition to focusing on basic processing schemes, this study also puts emphasis on practicality. Theoretical considerations are not used, but rather synthetic image flow fields are processed to determine expected algorithm performance characteristics in terms of measurement accuracy and computational efficiency. The flow fields selected span a range of particle displacements and flow types. In certain current techniques found in the literature, *a priori* knowledge of the flow is integral to the success of the processing method, making such methods flow-specific in many cases. In this investigation, no *a priori* knowledge of the flows is used in the processing, thus providing a truly comprehensive evaluation of the processing methods.

3.2 Processing digital image data for velocimetry

Traditional forms of DPIV processing algorithms have relied on auto- or cross-correlation schemes. Cross-correlation implementations typically are favored over auto-correlation methods since the former yield directionally unambiguous displacements and have been found to provide superior aggregate performance—*e.g.*, lower correlation noise, greater dynamic range, and less gradient bias (Keane and Adrian [67]). Early DPIV cross-correlation analyses were performed in the Fourier domain, as this provided the most efficient means of processing the data (*e.g.*, Willert and Gharib [144]). However, with the continual development of faster CPU engines, the need for the transformation into frequency space is becoming less imperative, and some methods of correlation are performed “brute force” in the spatial domain directly (Huang *et al.* [58]; Roesgen and Totaro [112]). Here, four cross-correlation algorithms are presented: (1) direct spatial domain correlation, (2) standard FFT (Fast Fourier Transform) correlation, (3) dynamic FFT correlation, and (4) a new hybrid dynamic FFT/spatial correlation. In addition, a particle tracking technique that is based closely on the hybrid method of Cowen and Monismith [20] is presented. Each scheme was investigated for accuracy and efficiency for a number of prescribed displacement fields.

3.2.1 Cross-correlation algorithms for DPIV

Extracting velocity information from a pair of single-exposed DPIV images involves subdividing the images (hereafter, \mathcal{A} and \mathcal{B}) into smaller subimages and quantifying the average spatial shift of particles from subimages in \mathcal{A} to subimages in \mathcal{B} . Determining the pixel offset between \mathcal{A} – \mathcal{B} subimage pairs (*i.e.*, the local flow displacement) can be accomplished using the statistical operation of cross-correlation. In two-dimensions, the discrete cross-correlation between two real-valued functions, $f(x, y)$ and $g(x, y)$, is given by Gonzales and Wintz [43] as,

$$\Phi_{fg}(x, y) = f \circ g = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n)g(x + m, y + n), \quad (3.1)$$

for $x = 0, 1, \dots, M - 1$ and $y = 0, 1, \dots, N - 1$. Equation 3.1 can be used to express the correlation between subimages f and g , of dimensions $A \times B$ and $C \times D$, respectively. Formally, these samples are assumed extended and periodic with periods M and N in the x and y directions such that $M \geq A + C - 1$ and $N \geq B + D - 1$; however, in practice this typically is not the case. The correlation can be computed in either the spatial domain directly or in the Fourier domain (with certain potential restrictions), and as mentioned above, the latter is preferred when computational efficiency is the primary goal. Before addressing issues of implementation, equation 3.1 should be slightly modified to yield a de-meaned, variance-normalized cross-correlation coefficient:

$$C_{fg}(x, y) = \frac{[f(x, y) - \bar{f}] \circ [g(x, y) - \bar{g}]}{\sqrt{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} [f(m, n) - \bar{f}]^2 [g(m, n) - \bar{g}]^2}}. \quad (3.2)$$

The quantities \bar{f} and \bar{g} represent the respective means of subimages f and g . A variance-normalized coefficient is a more robust measure of correlation between two subimages since it accounts for intensity variation (both local and global) between images \mathcal{A} and \mathcal{B} . This measure of correlation has been shown to be more accurate than alternative non-normalized measures by Burt *et al.* [15] and Huang *et al.* [57], for example. In each of the algorithms examined here, correlation peak locations are found to sub-pixel accuracy using two 3-point Gaussian curve-fits—one in the x -direction, one in the y -direction. Thus, in this investigation equation 3.2 was used as the correlation measure, leaving only the implementation of

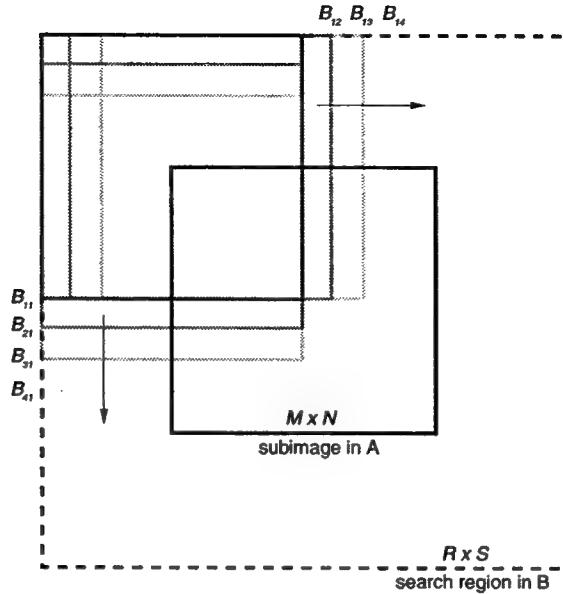


Figure 3-1: Schematic of direct spatial cross-correlation calculation. The shaded subimages (B_{ij}) show a sampling of all possible image \mathcal{B} subimages used to generate the complete correlation surface. The remainder \mathcal{A} - \mathcal{B} pairings are found by tiling the B_{ij} subimages horizontally and vertically within the search region as indicated by the arrows.

the correlation calculation (spatial domain versus frequency domain) and advanced subimage matching strategies as the differences between the DPIV algorithms considered.

Direct spatial correlation

While the direct spatial correlation is the most intuitively straightforward manner in which the cross-correlation can be computed, in that it mimics a simple template matching operation, it is also the most numerically intensive. Specifically, it uses the coefficient given by equation 3.2 with a direct multiplication operation like that given by equation 3.1 with the caveat that rather than imposing periodicity through extension, the actual content of image \mathcal{B} can be used. In this work, the spatial domain correlation has been implemented as follows. In figure 3-1, a subimage in \mathcal{A} of size $M \times N$ is identified, and a search region in \mathcal{B} of size $R \times S$ is centered about the subimage in \mathcal{A} . The subimage in \mathcal{A} is correlated with each possible $M \times N$ subimage in the search region of \mathcal{B} , building up a correlation surface of size $(R - M + 1) \times (S - N + 1)$. The location of the peak of this correlation surface indicates

the best estimate of the mean displacement of particles between the subimage pair; this holds true regardless of the scheme used to cross-correlate.

There are a number of advantages to direct spatial correlation. Most importantly, this approach minimizes the out-of-pattern (or, in-plane loss-of-pairs) effect (Keane and Adrian [67]; Huang *et al.* [58]) since, under ideal conditions, a properly defined search region allows for all of the particles in the subimage from \mathcal{A} to be found in a subimage from \mathcal{B} . This has been referred to as Particle Image Pattern Matching, or PIPM (Huang *et al.* [58]). As a result, direct spatial correlation can be considered a highly accurate method of cross-correlation, and this indeed was found to be the case as will be shown in section 3.3. A further advantage of the spatial domain approach is its flexibility. Unlike FFT-based schemes that can be restricted by certain Fourier transform algorithms (*e.g.*, square, power-of-2, equal-size subimages), the spatial correlation has no limitations on subimage/search region shapes or sizes. This can be advantageous when flows are quasi-unidirectional (*e.g.*, wall boundary layers). In such instances, the search region can be stretched in the flow direction and reduced in the cross-flow direction, reducing the number of computations necessary. Eliminating unnecessary correlations is particularly desirable for the spatial correlation. For an $N \times N$ subimage with a search region of dimensions $2N \times 2N$, the present spatial correlation requires $(N + 1) \times (N + 1)$ calculations, each of $\mathcal{O}(N^2)$ to build up the correlation surface. The direct spatial correlation is therefore an $\mathcal{O}(N^4)$ operation. To ameliorate this drawback, possible speed enhancements have been implemented for direct spatial correlation schemes (Roesgen and Totaro [112]; the present hybrid technique).

FFT correlation

Because the direct spatial cross-correlation operation is so computationally demanding, use of Fourier transform methods are popular in DPIV processing. In frequency space, the double summation in equation 3.1 can be replaced by a complex conjugate pointwise multiplication of the two-dimensional Fourier transformed subimages. This is expressed by the correlation theorem,

$$f(x, y) \circ g(x, y) \Leftrightarrow F(\xi, \eta)G^*(\xi, \eta), \quad (3.3)$$

where F denotes the Fourier transform of the function f , and G^* represents the complex conjugate of the Fourier transform of function g . As an example, the FFT implementation involves taking an $N \times N$ subimage from image \mathcal{A} and a *co-located* $N \times N$ subimage in \mathcal{B} , and generating the correlation surface in a single step using equation 3.3 with equation 3.2.

When working with an FFT calculation of the correlation, the issue of periodicity of the samples arises. Windowing methods like those used in spectral analysis to deal with the artifacts of non-periodic data are undesirable in DPIV since they can introduce systematic errors and/or degrade the correlation signal-to-noise ratio (Raffel *et al.* [108]). Furthermore, Westerweel [139] has argued that the use of window functions is unnecessary in practice. The method of zero-padding the subimage data is also a topic of debate. For a common $N \times N$ subimage, this typically would necessitate padding the sample with zeros to arrive at a sample of size $M \times M$, with $M \geq 2N$. Westerweel [139] notes that in practice, for particle displacements less than about 1/4 the subimage size ($N/4$), the cross-correlation effectively vanishes beyond displacement offsets of $N/4$. This being the case, the choice of $M = N$ is acceptable. However, calculating the cross-correlation without enforcing periodicity (without zero-padding) can have consequences for larger displacements. Specifically, Westerweel [139] finds that for displacements near $N/2$, without zero-padding, the variance of the inverted Fourier transform is increased by an amount proportional to the noise in the estimated cross-correlation. Neither zero-padding nor windowing is used in the FFT-based methods studied here.

The FFT correlation method, while only being $\mathcal{O}(N \log_2 N)$ in complexity, suffers dramatically from the out-of-pattern effect since the subimage in \mathcal{B} typically will fail to contain all of the particles found in the subimage from \mathcal{A} . As the mean particle displacement increases (*e.g.*, > 1 pixel) this effect becomes more pronounced. Note, the FFT method of correlation studied here is nearly identical to the normalized cross-correlation (NCC) method of Huang *et al.* [57] except that all values of the correlation are normalized in the present scheme whereas Huang *et al.* normalize only the values comprising the peak, sparing some computational effort. Also, the present coefficient uses de-means samples, and the Huang *et al.* NCC method does not.

Dynamic FFT correlation

Westerweel *et al.* [141] have shown that out-of-pattern effect difficulties can be removed by using an image shifting technique. The authors find that by shifting the location of the subimage in image \mathcal{B} by an amount equal to the integer part of the displacement (in pixels, px), the precision of the cross-correlation processing can be enhanced. The improved precision stems from results that have shown the root-mean-square (RMS) error in estimating displacements is proportional to the displacement itself for displacements less than 1/2 px (Adrian [3]; Willert and Gharib [144]). By shifting the location of the subimage in \mathcal{B} by the integer value closest to the total displacement and performing the cross-correlation, the residual displacement will always be less than 1/2 px, thus yielding less variation in the measurement error. The Westerweel *et al.* approach is useful when there exists a persistent mean component to the flow being studied, since the shift is done based on *a priori* knowledge of the mean displacement.

In their hybrid particle tracking algorithm, Cowen and Monismith [20] present a *dynamic* image shifting procedure for FFT-based correlation methods. This strategy was adopted for the dynamic FFT (DFFT) algorithm used here. The scheme begins with a standard FFT pass through the entire image. This pass provides an estimated mean displacement vector for each node location. For each node, the (dx, dy) displacement estimates are rounded to their nearest integer values, and the results are used as the amounts by which to dynamically shift the image \mathcal{B} subimage in the next pass. This process is repeated iteratively for all nodes on a node-by-node basis, stopping when the rounded integer-pixel displacements for both components of the displacement are zero. Five dynamic iterations are performed in all. For nodes where the displacement fails to converge to zero or where the dynamically located subimage in \mathcal{B} forces the subimage beyond the image domain, the displacement vector is determined via nearest neighbor bilinear interpolation. A similar DFFT method is reported by Scarano and Riethmuller [116] in which, along with dynamic shifting, iterative subimage size refinement takes place, allowing for increased spatial resolution.

Hybrid correlation

To achieve both high accuracy and efficiency, a *hybrid* cross-correlation scheme was developed in this thesis, which makes use of the dynamic FFT approach and the direct spatial calculation of the correlation. It is a straightforward extension of both that exhibits the favorable accuracy found in the direct spatial correlation calculation along with the superior computational efficiency of the FFT methods. The first phase of the hybrid scheme performs a full DFFT correlation operation for the image pair. The second phase further refines the cross-correlation peak with a limited spatial domain correlation operation. An example serves to illustrate the methodology. Consider a node located at (x, y) that has yielded a converged DFFT displacement result of $(-1.017 \text{ px}, 6.230 \text{ px})$ using $N \times N$ subimages. The next step locates an $(N + 6) \times (N + 6)$ search region in image \mathcal{B} centered at coordinates $(x - 1, y + 6)$. This yields a 7×7 correlation surface that is dominated by the correlation peak. The choice of a 7×7 correlation surface is somewhat arbitrary, but it allows for refinements to the DFFT displacement up to 2.5 px based on the operation of the present sub-pixel peak locator, with an acceptable increase in computation time. This refinement pass is done once for all nodes in the image.

3.2.2 Particle tracking

Spatial resolution is a limitation in all correlation-based DPIV approaches. Flow scales smaller than the subimage dimensions cannot be captured since the result from an auto- or cross-correlation operation represents the mean displacement of all particles within the subimage region. Therefore, subimage size governs spatial resolution of the flow. Reducing the size of the subimage increases the resolution but there is a limit to how small the subimage can be based on the statistical need for an adequate number of particle images, or unique image texture, in the subimage. A further limitation of typical correlation procedures is their difficulty in handling velocity gradients. This also derives from the integrating nature of the correlation operation. The main problem encountered when velocity gradients are present (as is almost always the case for fluid flows of interest), particularly for standard static FFT correlation schemes, is gradient biasing. For instance, when an in-plane gradient exists, faster moving particles leave the subimage region and the remaining, slower moving

particles act to bias the displacement estimate toward slower velocities. This error is reduced by using dynamic methods like the ones described here. However, even with dynamic correlations that achieve PIPM, small scale gradients are still undetectable because of the averaging done by the correlation. An alternative technique that can overcome resolution and gradient biasing effects is particle tracking. Particle tracking, by definition, tracks individual flow tracer particles and, assuming ideal particles that exactly follow the flow, can provide a high accuracy measurement of the flow field. The challenge is to successfully track large numbers of particles to achieve high spatial resolution.

In this thesis work, a particle tracking scheme based on the ideas of Cowen and Monismith [20] was implemented. The scheme is also similar to those of Keane *et al.* [68] and Rehm and Clemens [111]. It is a hybrid technique that uses dynamic FFT DPIV results to guide the particle matching algorithm. This approach allows for much higher particle seeding densities than have been used in previous particle tracking methods. Details of the method can be found in Cowen and Monismith [20], therefore only the salient aspects are highlighted. The routine begins by performing a coarse dynamic FFT DPIV pass over the entire image with non-overlapping subimage regions. Non-converged nodes are interpolated, and the entire displacement field is processed through an outlier detection operation (see page 66 for details). Next, the original images are binarized using a user-adjustable threshold, and individual particles in both images are identified and labeled using a region-growing algorithm. Particle matching then takes place either from image \mathcal{A} to image \mathcal{B} (number of particles in \mathcal{A} > number of particles in \mathcal{B}), or from \mathcal{B} to \mathcal{A} (number of particles in \mathcal{B} > number of particles in \mathcal{A}) and the results reversed. To find a particle match, several steps are taken. First, an initial displacement estimate is determined. The present code has two options for this step. The crudest and most efficient option considers the location of the particle to be matched, and based on that, the local DPIV result is used as the initial displacement estimate. The more sophisticated and time-consuming option generates a global displacement estimating function using a Hardy multiquadratic interpolation scheme (Hardy [49]) as in the Cowen and Monismith version. In this approach, each particle is assigned a unique displacement estimate. When the two methods were compared, the results of the particle tracking using the interpolation scheme were not improved significantly,

if at all, and generally did not warrant the added computational time necessary to solve the system of equations involved. In the next step, a final refinement FFT cross-correlation operation between an 8×8 subimage centered on the particle to be matched in image \mathcal{A} and a subimage of the same size displaced in image \mathcal{B} by the DPIV estimate is performed. The resulting peak correlation coefficient is compared to a user-selectable correlation coefficient threshold. If the computed value fails to meet the threshold, the particle is not tracked further. This avoids match errors due to poorly correlated local regions of the image. Any adjustments to the initial estimate as a consequence of this operation are included and the result is the final guess displacement. Lastly, a 3×3 subimage is located in the second image based on the final estimated displacement, and if a particle is found in this box it is determined to be the matching particle. The displacement is found by differencing the two particle centers, which are found from either an intensity-weighted centroiding method or, as in the Cowen and Monismith approach, two 3-point Gaussian fits. A final quality control check is imposed. The particle displacement, dp , is compared to the quantities $(1 \pm \alpha)U_{guess}$, where U_{guess} is the final guess displacement and α is a selectable factor. If $(1 - \alpha)U_{guess} < dp < (1 + \alpha)U_{guess}$, the result is retained, otherwise, the particle match is discarded. This check effectively operates like an outlier removal routine and is successful at removing wildly incorrect particle matches. All particles are tracked in the same fashion, and the resulting irregularly spaced data are put to a regular grid using a standard triangulation-based interpolation routine.

3.3 Accuracy and efficiency

To ascertain the accuracy and computational efficiency of the five processing schemes described above, numerically generated DPIV images were used with known prescribed displacements. Images were generated by populating a black background with randomly located particle images. A random number generator was used to determine the horizontal and vertical coordinates of each particle image in the field. Particle images were assumed

to possess normally distributed intensity profiles of the form

$$I(x, y) = I_o \exp \left[\frac{(x - x_o)^2 + (y - y_o)^2}{-2\sigma^2} \right], \quad (3.4)$$

where I_o is the central intensity (also randomly generated) at the particle center (x_o, y_o) , and σ is related to the particle diameter as $d_p = 4\sigma$ (based on the e^{-2} intensity value). A constant particle diameter of $d_p = 2.8$ px was used throughout. The digital image array was filled by integrating equation 3.4 over each member pixel for each particle image. An 8-bit digitization operation was assumed, and pixel values were truncated to 255 when necessary. The discretization of particle images has the effect of introducing additive white noise to the synthetic images (Westerweel [140]). The particle density was set at 20 particles per 32×32 image region (0.02 particles/px²). This density is realistic in practice and is comparable to values used in other simulated image tests. Once particles were distributed in image \mathcal{A} , a displacement function was used to shift the particle locations in image \mathcal{B} and the pixel intensity integrations repeated. Three displacement functions were investigated: (1) uniform translation in a single direction, (2) uniform solid body rotation, and (3) a group of diverging two-dimensional stagnation flows. Out-of-plane particle motions were not simulated.

Errors associated with DPIV algorithms often are decomposed into two components: mean bias error and RMS error. Given the actual displacement, d_a , and the estimated DPIV/PTV displacements, d_i ($i = 1, 2, \dots, N$), the mean bias error can be defined as

$$E_{bias} = d_m - d_a, \quad (3.5)$$

where d_m represents the spatial mean of the N estimated displacements. The RMS error was defined as

$$E_{rms} = \sqrt{\frac{1}{N} \sum_{i=1}^N (d_i - d_m)^2}, \quad (3.6)$$

and reflects the deviation of the displacements about their mean. Mean bias errors arise when the actual correlation peak is not well represented by the peak fitting curve or procedure (Huang *et al.* [57]). This is often the case when out-of-pattern effects are prominent, or

when strong gradients within the subimage act to broaden and/or distort the symmetry of the correlation peak. RMS errors are due to any number of influences: poor particle seeding, velocity gradients, out-of-plane particle motions, imaging imperfections (*e.g.*, non-uniform illumination, camera non-linearities, parallax), and data acquisition noise. As in Huang *et al.* [57], the errors associated with the methods of processing the digital image data are the main focus. For all data shown here, the only post processing done on the displacement fields was a simple outlier removal operation. Outliers were determined on a node-by-node basis by differencing each component of displacement with its eight neighbors and comparing the results to a user-selected threshold. If the differences exceeded this threshold for more than four neighbors, the value was flagged an outlier, and replaced with a bilinearly interpolated value. With slight modification to the above scheme, edge and corner values in the data arrays were processed also with the outlier removal operation. Use of the outlier removal operation was consistent with the theme of practical application of these methods.

It should be noted that the error in estimated displacements and the generation of spurious displacement estimates are decoupled (Westerweel [139]). The absolute accuracy (or equivalently, error) stems from, amongst other factors, the ability to accurately estimate the correlation peak center, which depends on the width and shape of the peak. For a uniform displacement, the peak width is determined by the particle image diameters. Too large particles will act to broaden and degrade the peak, which ideally should resemble a Dirac delta function. Spurious displacement estimates, on the other hand, arise when a random correlation peak is identified as the correct displacement peak. The probability of this occurrence is inversely proportional to the particle image density. Therefore, increasing the particle seeding density does not add to the accuracy of the measurement, but rather it increases the probability of a correct measurement. Ultimately, the performance of DPIV processing is a combination of these two factors: identifying the correct displacement peak and determining the location of the peak center accurately.

3.3.1 Uniform translation

Following the approach of Huang *et al.* [57], a series of sub-pixel uniform displacements ranging from 0 to 1 px was considered first. Image size was 640 px \times 640 px. The situations

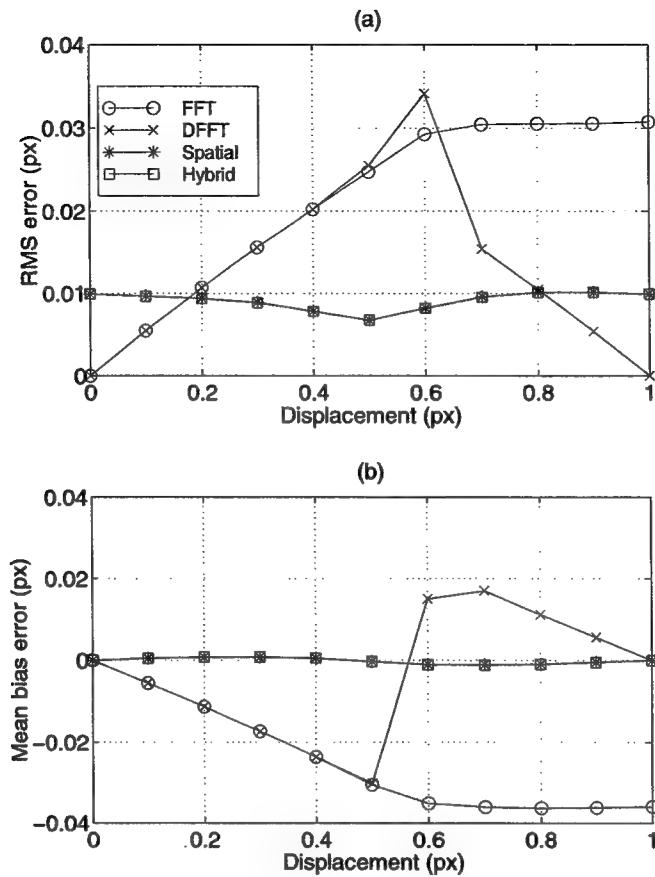


Figure 3-2: Variation of (a) RMS error, and (b) mean bias error, for all four DPIV methods over the 0–1 px displacement range with 32×32 subimage processing.

for 16×16 , 32×32 , and 64×64 subimage size processing were studied separately with zero overlap (all measurements independent). For each case, the processing strategy was held consistent for all processing methods (*e.g.*, search region dimensions were set at twice the subimage dimensions in all cases). In these tests, outlier removal was not necessary. The PTV scheme used a correlation threshold of 0.85 and an α -factor of 0.5 throughout. For all cases, the PTV method used the simple method of generating the initial displacement guesses.

The DPIV processing results for the RMS and mean bias errors are shown in figure 3-2 for the 32×32 case only. Because direct spatial correlation incorporates the most image information (yielding robust correlation peak data) and does not suffer from out-of-pattern

effects, this method is expected to produce the most accurate results, and this is observed to be the case. This figure also begins to demonstrate the favorable performance of the hybrid scheme and the reasoning behind it. The hybrid technique is able to retain the same accuracy as the full spatial correlation with an increase in speed of at least an order of magnitude (as will be shown in section 3.3.2).

Considering the RMS errors, the variability of the error for the spatial and hybrid techniques remains essentially constant over the complete sub-pixel range. The reason for this finite variability is the statistical variability in the content of the individual subimages themselves. With direct spatial calculation, the peak is comprised of correlations between several different A - B subimage combinations, and consequently the peak can assume an asymmetric shape due to the particular spatial distributions of the particle images and their intensities in these subimages. This asymmetry should be a random process assuming randomly located particles with randomly generated intensities. Thus, sub-pixel Gaussian curve-fitting of the peak, as was done here, will yield random variability in the computed peak location. FFT-based methods produce less correlation peak asymmetry near zero displacements. This is observed for the FFT and DFFT methods near displacements of zero, and for the DFFT near 1 px, where it has shifted to produce a zero residual displacement. The standard FFT correlation shows a proportional increase in RMS error as the pixel displacement increases for displacements less than 1/2 px, confirming what was noted earlier. Beyond 1/2 px, but less than 1 px, the RMS error is nearly independent of the displacement. The behavior of the DFFT scheme is somewhat more subtle. For displacements less than 1/2 px, the DFFT results are identical to that of the FFT method since no shifting takes place. At 0.7 px and beyond, the DFFT results are symmetric (because of shifting), and at 0.5 px and 0.6 px there are discrepancies. The explanation lies with the dynamic nature of the DFFT method. When the displacement is 1/2 px, most DFFT results will converge to rounded values of zero px because of the typical underestimation of displacements computed using FFTs (see figure 3-2b). A small number of nodes may converge to 1 px, and for these nodes the DFFT then will shift the second subimage by 1 pixel. Consequently, the mean result at 1/2 px will be different for the DFFT compared to the FFT, which does not shift. Similarly, at 0.6 px most DFFT nodes will converge to 1 px, but a small number may round

to zero px because of underestimation. This results in the increased variability seen at this displacement.

Turning to the mean bias errors, for the spatial and hybrid methods, the mean bias incurred is extremely small (0.0006 px mean magnitude over the complete 0–1 px range). The direction of the mean bias for these two methods indicates a trivial degree of inverse “peak-locking” as discussed by Fincham and Spedding [32]. Peak-locking is an effect attributed to fitting asymmetric peak data with symmetric curves such as Gaussians and parabolics. For the two FFT-based methods, the mean bias errors are comparable in magnitude to the RMS errors, a result also described by Huang *et al.* [57]. The non-symmetric behavior of the DFFT mean bias errors over the 0–1 px range follows the same logic as described above for the RMS errors. Evident is the consistent underestimation of the displacement using the static FFT scheme. This is a consequence of the skewness of the correlation peak that biases the estimated displacement toward zero. This skewness can be thought of as deriving from a small, yet finite, amount of particle image content “leaking” out of (decreasing the correlation in the true direction of displacement) and “seeping” into (increasing the correlation in the opposite direction) the subimage in \mathcal{B} . Westerweel [139] has shown that this bias is directly proportional to the width of the correlation peak relative to the subimage dimension. In general, this width is determined by not only the particle image size but also by the variation of the displacement field over the subimage. To further amplify this behavior, figure 3-3 shows the displacement frequency distributions due to a 1.0 px shift for the FFT and direct spatial correlation methods. This result clearly shows both the significant negative mean bias error and the substantial variability of the displacement estimates from the FFT calculation. Westerweel [139] has shown an analytical result that enables elimination of this mean bias error for FFT methods (also, Raffel *et al.* [108]). The FFT implementation studied here does not incorporate this correction. Essentially, the correction involves a compensation factor in the correlation calculation to account for in-plane loss-of-pairs. Interestingly, the present hybrid technique achieves the same effect in a different manner: the FFT implementation of the correlation is employed and the effect of in-plane loss-of-pairs is eliminated via the local, direct spatial calculation of the correlation peak.

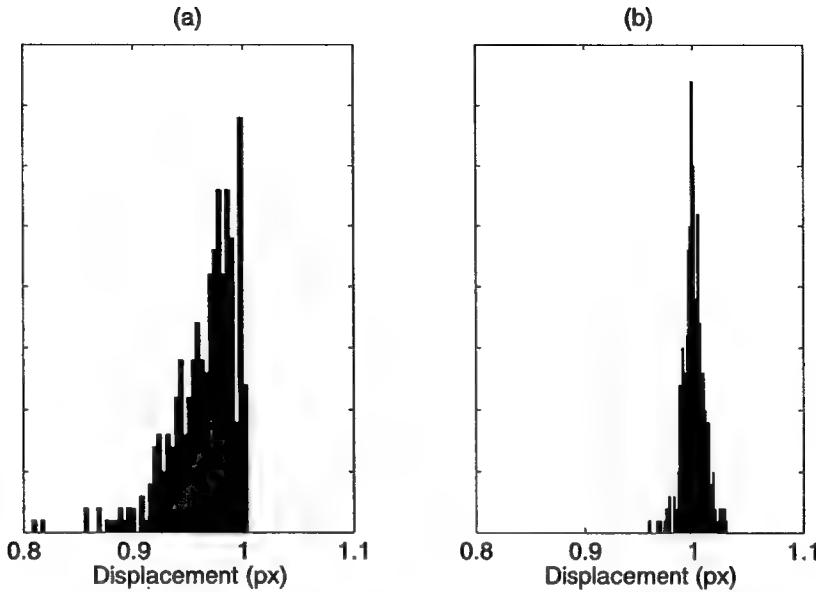


Figure 3-3: Displacement frequency distributions for (a) FFT and (b) direct spatial correlation methods. Actual displacement is 1.0 px.

Comparing the results of figure 3-2 to that of Huang *et al.* [57], there exist certain differences that warrant clarification. Their errors were found to be symmetric about the 1/2 px mark, while those here are not symmetric for the FFT method, and only symmetric for the DFFT method away from 1/2 px displacement. It is believed that this difference is due to the fact that in the present tests, the FFT scheme remained static, and the DFFT scheme was allowed to shift according to its rules of rounding, while in the Huang *et al.* work, the shifting presumably was forced ideally based on the known displacements. While such shifting would avoid the deviations from symmetry found in the present results, it would also defeat the purpose of a dynamic processing scheme, and it was felt that allowing the DFFT algorithm to proceed unaided was in keeping with the practical implementation of DPIV analysis this study addresses.

Processing with subimages of sizes 16×16 and 64×64 yielded similar qualitative behaviors in both components of the error over the 0–1 px range. Summarizing, for the 16×16 case, RMS errors were higher by about a factor of 4 and mean bias errors were higher by about a factor of 2.5 over that shown for the 32×32 result. For the 64×64 case, RMS errors

were lower by about a factor of 3 and mean bias errors were lower by about a factor of 2 compared to the 32×32 result. This implies that the error in the displacement estimates decreases with increasing subimage size, and this would be expected since larger subimage sizes incorporate a greater statistical sample, or alternatively, more unique subimage templates.

Investigating the errors for pixel displacements in the 0–1 px range is useful in ascertaining the detailed behavior of the correlation schemes. Here, it has revealed certain shortcomings of FFT, as well as DFFT, methods of DPIV processing that are undesirable in certain flow situations. However, rare is it that DPIV images yield such small displacements in general. Huang *et al.* [57] note that displacements larger than 1 pixel can be corrected for with the image shifting technique (Westerweel *et al.* [141]), but this is only the case when the necessary shift is known, and applies to the entire image (*e.g.*, a channel flow). To examine the behavior of these processing methods for pixel displacements larger than 1 pixel without artificial shifting, we performed the same procedure as above for displacements of 1.25, 2.25, ..., 10.25 px (figure 3-4). First to note is the relatively poor performance of the standard FFT method, in terms of consistent negative mean bias error (underestimation) and larger RMS variability in the errors. This is a consequence of its static nature and the resulting out-of-pattern effects. To illustrate the underestimation effect of the FFT method in practice, figure 3-5 is included, which shows a laboratory result from a flume calibration experiment. The goal of this experiment was to calibrate the flow field in a newly constructed flume, and when an FFT correlation technique was used for processing, the results were underestimated based on comparison to the spatial correlation results. While the discrepancy is not large, the fact remains that the FFT method is less accurate.

Returning to figure 3-4, for the other three methods, patterns very similar to the results for the 0–1 px range are repeated for subsequent single pixel intervals (see figure 3-9), and the results at $dx.25$ px are what are observed. Thus, increasing displacements have little effect on the behavior of the accuracy of these three methods. This is noteworthy because image shifting is taking place automatically, thereby handling general flow patterns without any special knowledge of the flow. Results using 16×16 subimages showed both

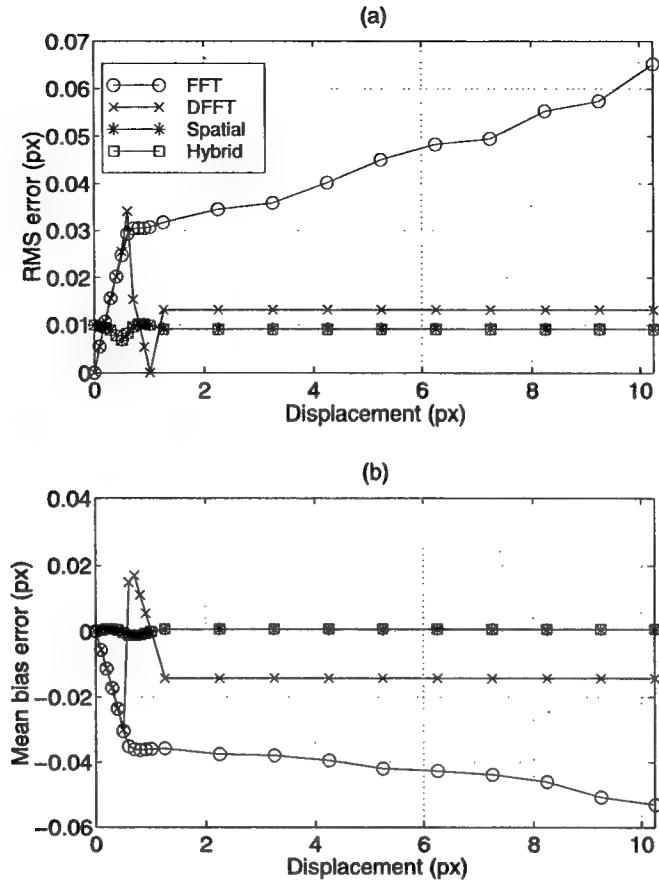


Figure 3-4: Variation of (a) RMS error, and (b) mean bias error, for all four DPIV methods over the 0–10.25 px displacement range with 32×32 subimage processing.

larger RMS and mean bias errors, and somewhat less systematic behavior over the range of displacements. Considering just the FFT results as an indicator, the maximum RMS error was nearly an order of magnitude higher (at a displacement of 5.25 px—beyond this point, the correlation is nearing the Nyquist limit and the estimates become poor) compared to the 32×32 case. The largest FFT mean bias error was higher by roughly a factor of 3 for the 16×16 case. Both error components for the FFT method were reduced when processing with 64×64 subimages: the peak RMS error was reduced by approximately a factor of 4, and the peak mean bias error was reduced by about a factor of 2.5.

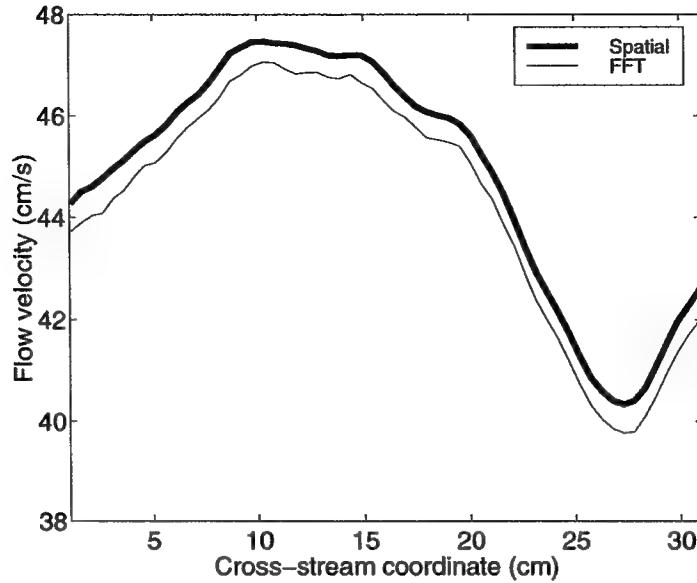


Figure 3-5: Flume calibration experiment result. Shown is the streamwise component of the flow taken along the cross-stream direction (cross-channel).

The mean, relative displacement error,

$$E_{rel} = \frac{1}{N} \sum_{i=1}^N \left| \frac{d_i - d_a}{d_a} \right| \times 100\%, \quad (3.7)$$

for all subimage size cases over the detailed 0–1 px range and the extended displacement range is shown in figure 3-6 (16×16), figure 3-7 (32×32), and figure 3-8 (64×64). In each case, the spatial and hybrid methods show the smallest errors, with the DFFT method yielding similar results. For all methods, the errors beyond about 2 px are relatively constant—less than 5% for the 16×16 case, less than 2% for the 32×32 case, and less than 1% for the 64×64 case. These figures show that the directly calculated correlation methods (spatial and hybrid) are noticeably superior to the traditional static FFT method over the complete displacement range studied, and are slightly better than the DFFT method over this range.

To further understand the detailed behavior of the DPIV errors at larger displacements, a more detailed analysis was performed over the 5–6 px displacement range, and the 32×32 results are given in figure 3-9. The errors for the DFFT, spatial, and hybrid methods appear

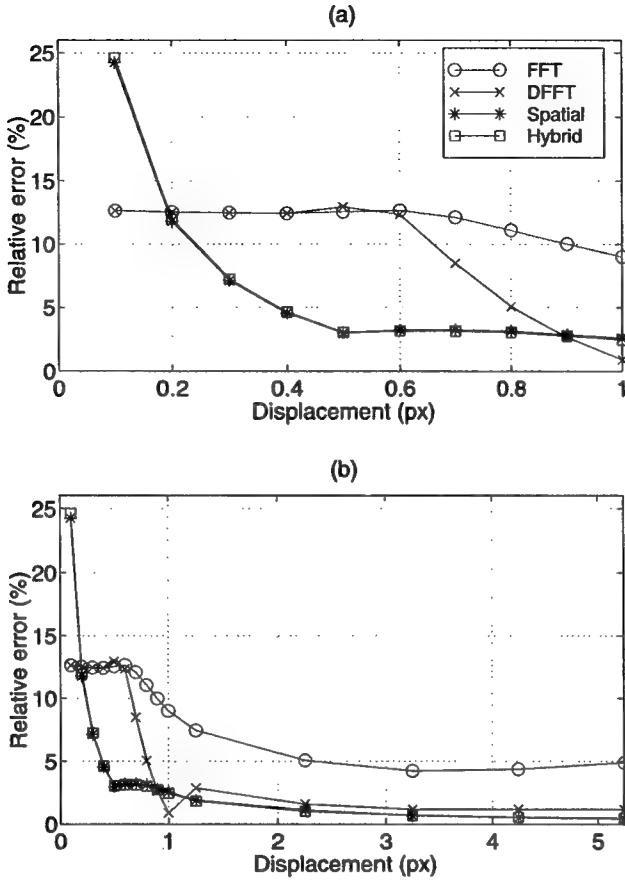


Figure 3-6: Variation of relative error over the (a) 0–1 px displacement range, and (b) 0–5.25 px displacement range, for all four DPIV methods with 16×16 subimage processing.

to be cyclical over the displacement range studied. Over this range, both the RMS and mean bias errors show qualitative similarity to the 0–1 px range results, indicating that PIPM is achieved. This is not the case for the FFT method, and consistent errors are observed. Unlike the result for the 0–1 px range, the relative error also exhibits a cyclic pattern over the 5–6 px displacement range. This is because the relative error for the 0–1 px range is based on scaling the absolute error by a small displacement value (< 1 px), whereas for the 5–6 px range, the absolute error is similar but the scaling value is much larger.

At this point, the results of PTV processing of the uniform displacement tests are included. Figure 3-10 and figure 3-11 show the 32×32 DPIV error results with the raw, irregularly gridded results of PTV included. Although exhibiting favorable mean bias

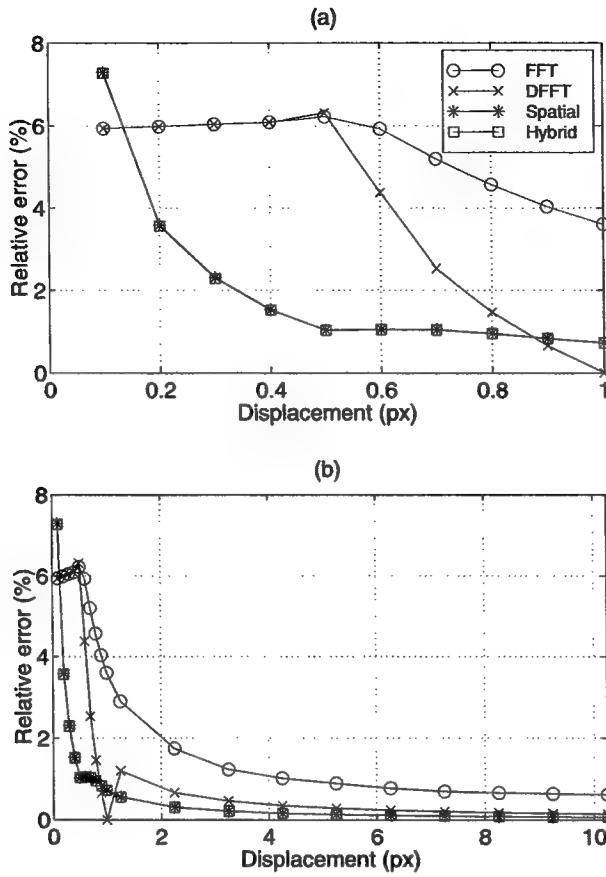


Figure 3-7: Variation of relative error over the (a) 0–1 px displacement range, and (b) 0–10.25 px displacement range, for all four DPIV methods with 32×32 subimage processing.

errors of less than 0.012 px, the RMS errors are quite large. The primary reason behind this is particle pair mismatching. Although these events were rare, the errors incurred were relatively significant, leading to the large RMS errors seen in figure 3-10a, and the large relative errors for the smallest displacements in figure 3-11a. Figure 3-11c shows the behavior of the PTV error at larger displacements in more detail. The PTV error also appears to exhibit a cyclic structure for displacements greater than 1 px.

3.3.2 Uniform rotation

To simulate rotational flow, synthetic image pairs were used to achieve constant vorticity, solid body rotations. Three rotations were investigated. The rotation used for the 16×16

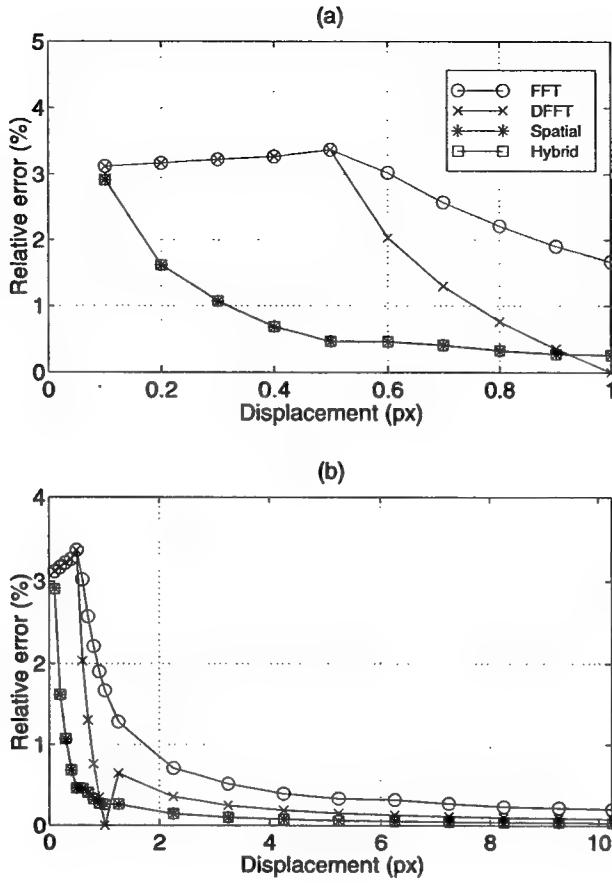


Figure 3-8: Variation of relative error over the (a) 0–1 px displacement range, and (b) 0–10.25 px displacement range, for all four DPIV methods with 64×64 subimage processing.

subimage case had a maximum pixel displacement of 5.6 px, the rotation for the 32×32 case had a maximum of 11.3 px, and the 64×64 case, a maximum of 22.5 px. The choice of these displacements was guided by the often used PIV rule-of-thumb that states that the maximum displacement within a given window should be less than about 1/3 of the window dimension (*e.g.*, Willert and Gharib [144]). In addition to assessing accuracy, these tests also were used as the basis for the algorithm efficiency comparison. Computational performance is presented in table 3.1. Computational times are the result of processing with a 400 MHz Intel PC running Linux. The speed of the standard FFT method is readily apparent; it is at least twice as fast as the other DPIV schemes, being 1–2 orders of magnitude faster than the direct spatial calculation. While not as efficient as the FFT-based algorithms,

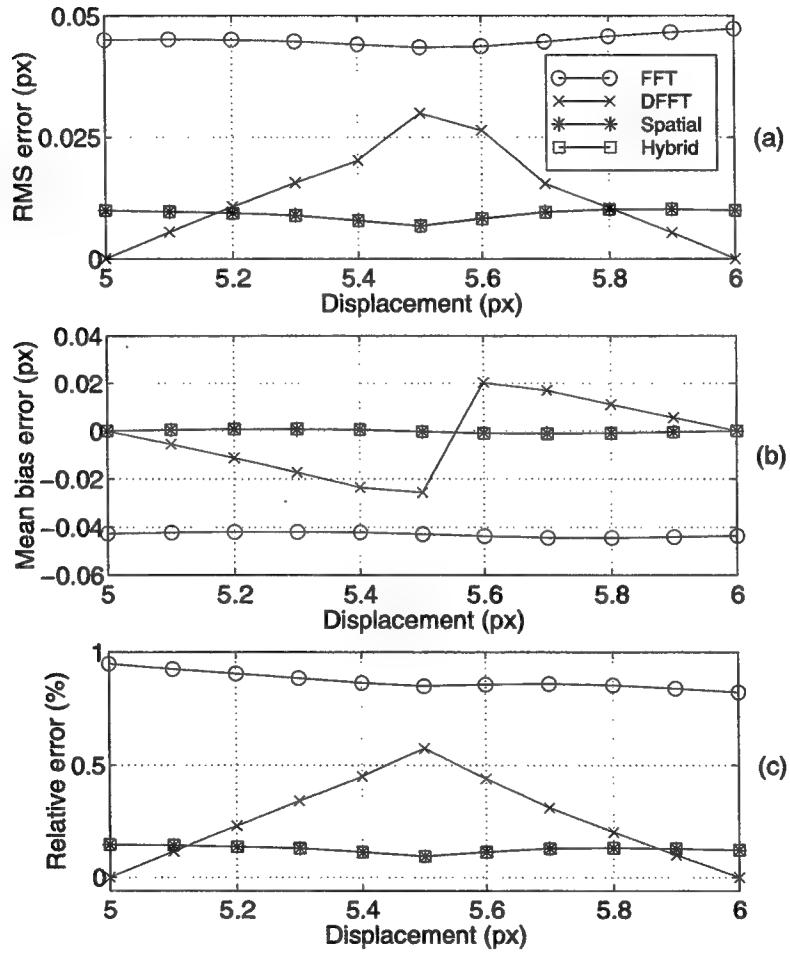


Figure 3-9: Variation of (a) RMS error, (b) mean bias error, and (c) relative error, for all four DPIV processing methods over the 5–6 px displacement range with 32×32 subimages.

(a) 16×16 subimages

Scheme	Time (sec)	Vector count	Vectors/sec
FFT	0.80	1521	1901
DFFT	1.68	1521	905
Spatial	10.90	1521	140
Hybrid	3.53	1521	431
PTV	5.20	4980	958

(b) 32×32 subimages

Scheme	Time (sec)	Vector count	Vectors/sec
FFT	0.86	361	420
DFFT	1.73	361	209
Spatial	35.81	361	10
Hybrid	3.33	361	108
PTV	5.28	4969	941

(c) 64×64 subimages

Scheme	Time (sec)	Vector count	Vectors/sec
FFT	1.66	81	49
DFFT	3.71	81	22
Spatial	123.89	81	0.65
Hybrid	5.39	81	15
PTV	8.18	4826	590

Table 3.1: Computational performance of DPIV/PTV processing algorithms.

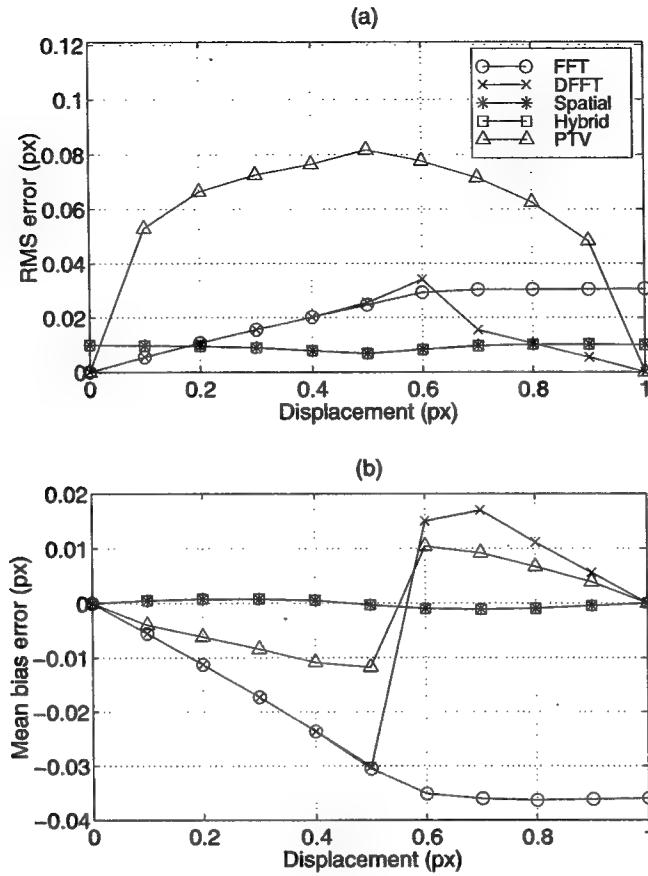


Figure 3-10: Same as figure 3-2 but with raw, ungridded PTV results included.

the hybrid technique is found to be anywhere from about 3 to about 20 times the speed of the direct spatial correlation. For the PTV results, subimage size reflects the size used during the coarse DPIV pass. Also, the PTV result for the 64×64 case made use of the more complex interpolation scheme to generate the initial displacement guesses. This added approximately half a second to the execution time over that of the simple scheme used in the other two cases. The speed of the PTV scheme is impressive considering the large number of vectors the method is able to yield.

The outcome of the accuracy tests is summarized in table 3.2. Outlier removal became a factor in the rotation tests, and is reported in this table. In this case, where a single mean

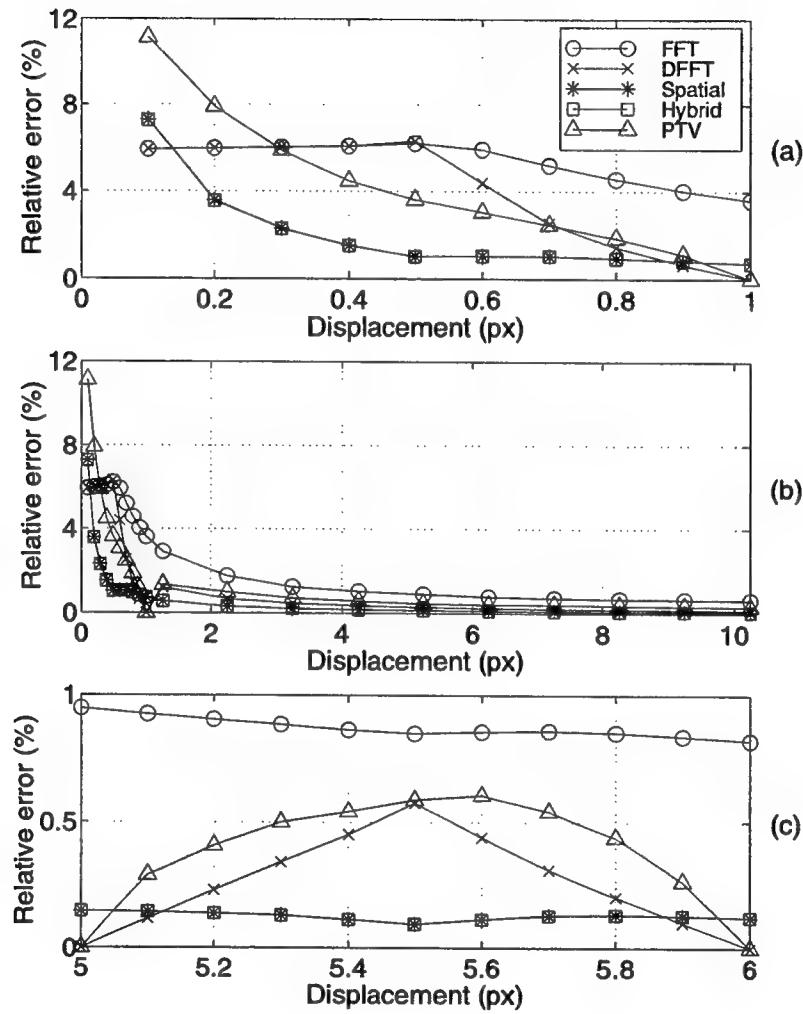


Figure 3-11: Variation of relative error over the (a) 0–1 px displacement range, (b) 0–10.25 px displacement range, and (c) 5–6 px displacement range, showing raw, ungridded PTV performance (32 × 32 subimage size used for all methods).

(a) 16×16 subimages

Scheme	Outlier fraction (%)	RMS error (px)	Relative error (%)	Vorticity error (%)
FFT	11.4	0.283	6.12	24.33
DFFT	4.6	0.137	2.75	13.67
Spatial	0.5	0.065	1.88	6.91
Hybrid	0	0.070	1.92	7.06
PTV	12.2	0.114	1.13	—
$\text{PTV}_{\text{DPIV-grid}}$	—	0.089	1.18	7.72
$\text{PTV}_{\text{fine-grid}}$	—	0.078	1.15	14.89

(b) 32×32 subimages

Scheme	Outlier fraction (%)	RMS error (px)	Relative error (%)	Vorticity error (%)
FFT	0	0.098	2.47	2.40
DFFT	0	0.062	1.97	2.19
Spatial	0	0.058	1.82	1.78
Hybrid	0	0.058	1.82	1.78
PTV	11.6	0.210	0.77	—
$\text{PTV}_{\text{DPIV-grid}}$	—	0.100	0.80	2.86
$\text{PTV}_{\text{fine-grid}}$	—	0.126	0.83	9.98

(c) 64×64 subimages

Scheme	Outlier fraction (%)	RMS error (px)	Relative error (%)	Vorticity error (%)
FFT	0	0.177	5.31	1.46
DFFT	0	0.188	5.31	1.48
Spatial	0	0.185	5.42	1.43
Hybrid	0	0.193	5.48	1.56
PTV	12.5	0.493	0.74	—
$\text{PTV}_{\text{DPIV-grid}}$	—	0.304	0.90	1.31
$\text{PTV}_{\text{fine-grid}}$	—	0.357	0.75	10.74

Table 3.2: Simulated uniform rotation results. Boldface indicates best performance. DPIV processing parameters are identical to those used in the uniform translation tests except for outlier detection. For the 16×16 case, an outlier threshold of 2 px was used; for the other two cases, a threshold of 4 px was used. DPIV outlier fraction represents the fraction of the total number of computed vectors that were determined to be outliers. PTV outlier fraction represents the fraction of the total number of particles that were eliminated from tracking due to poor local correlation or considerable deviation from the local mean displacement. In the case of PTV, a correlation threshold of 0.70 and an α -factor of 0.5 was used for all three cases. PTV for the 64×64 case used the multiquadratic interpolation scheme for the initial displacement guesses; the other two cases used the simple displacement estimate.

displacement is of little use, the definition of the RMS error was modified slightly,

$$E_{rms} = \sqrt{\frac{1}{N} \sum_{i=1}^N (d_i - d_a)^2}, \quad (3.8)$$

the actual local displacement being used instead. The mean (relative) vorticity error was defined as

$$E_\omega = \frac{1}{N} \sum_{i=1}^N \left| \frac{\omega_i - \omega_a}{\omega_a} \right| \times 100\%, \quad (3.9)$$

where $\omega_a = 2\theta$, θ being the angle of rotation. The estimated vorticities, ω_i , were computed from the local 8-point circulation calculation around each node. This approach has been shown to be preferable over a first-order central difference method (Westerweel [139]). Table 3.2 reveals how the performance of the algorithms is affected by the subimage size used. In terms of displacements, the 16×16 and the 64×64 subimage DPIV processing show the greatest error and variability in accuracy. For the 16×16 case, this is the effect of the reduced statistical sample size increasing the probability of spurious correlation peaks. Typically, at least four or five particles per subimage are required to obtain an unambiguous measurement of the displacement—at the particle seeding density used here, the average number of particles per 16×16 subimage is slightly greater than five. For the 64×64 case, the error is due to the spatial averaging of the larger subimage size—displacement gradients are being smoothed by the coarse spatial resolution.

Examining the algorithms, at the 16×16 subimage size, the FFT algorithm shows significant RMS error and large error in the derived vorticity. A dynamic FFT scheme reduces these errors by approximately a factor of 2. The direct spatial and hybrid methods further reduce this error by another factor of 2. Thus, at this fine spatial resolution, the spatial and hybrid techniques are still viable whereas the FFT-based routines are being strained. This would indicate that in addition to extending the usable valid dynamic range, the spatial and hybrid techniques also can improve upon the spatial resolution without sacrificing accuracy. The 32×32 subimage case results show much improved accuracy and less performance variability. All DPIV methods show favorable results, the differences in errors, both absolute and relative, being small. Note, however, that the spatial and hybrid schemes reduce the displacement errors by nearly a factor of 2 over that of standard FFT

processing. For the 64×64 case, it appears all methods perform equally. Based on these results, the errors of the DPIV methods would be expected to become less disparate as the subimage size became larger, at which point choice of a processing method can be done based on efficiency concerns alone.

Three forms of particle tracking results are also included in table 3.2. The first, PTV, represents the accuracy based on the irregularly spaced individual particle vectors. The second, $\text{PTV}_{\text{DPIV-grid}}$, represents the results after transforming the raw PTV vectors to a regular grid identical to the DPIV grid formed by the other four processing methods. The third, $\text{PTV}_{\text{fine-grid}}$, represents the raw particle tracking results gridded to a fine resolution regular array. The array density was chosen such that approximately the same number of vectors appear in the gridded result as in the raw result, thereby mimicking the actual spatial resolution.

In all three subimage size cases, the PTV result, which successfully tracked 80–85% of the detected particles, shows the smallest relative displacement error. Examining the 16×16 subimage case, the smoothing effect of the gridding procedure is observed, yielding a reduction in the particle tracking RMS errors. This smoothing also has the effect of marginally increasing the relative error as would be expected for a vortical flow with displacement gradients. The incurred vorticity error as calculated from the $\text{PTV}_{\text{DPIV-grid}}$ result is essentially the same as that for the two spatial correlation DPIV methods. For the fine-grid result, the vorticity error is doubled ($1.93 \times$). This is not surprising since the finer grid is close to twice ($1.82 \times$) the resolution of the DPIV grid. The vorticity calculation as implemented here involves a division by the node spacing and will show greater sensitivity as the node spacing is reduced. Thus, the errors in displacement, which essentially remain the same for both PTV gridded resolutions, are amplified during the vorticity calculation for the finer grid resolution. Similar PTV results are found for the 32×32 subimage case; however, the RMS errors for the raw data are nearing $1/4$ px. It should be kept in mind, though, that this flow field has significant pixel displacement; this fact is reflected by the excellent relative error results. It appears that the gridding operations in the 32×32 case have the same effect of reducing the RMS scatter and increasing the relative errors of the PTV results. The behavior of the relative vorticity error follows the earlier discussion.

In this case, the PTV_{fine-grid} results are at a spatial resolution 3.68 times the DPIV grid results, and as expected, the vorticity error increases by a similar factor of 3.49 since the displacement errors do not change significantly between the two resolutions. Similar behavior of derived quantities at higher resolutions also was reported by Rehm and Clemens [111]. The 64 × 64 PTV results reveal some interesting points of the technique. For this subimage size, the Hardy multiquadratic interpolation scheme used to determine the initial estimate displacements performed noticeably better than the simple method used for the other two subimage sizes. This is due to the large nature of the 64 × 64 subimage and the appreciable spatial averaging that takes place. By using a displacement estimate derived from a large region, a greater number of particles within that region can exhibit actual displacements that deviate considerably from this mean estimated displacement. Consequently, successfully tracking these particles is problematic. By employing the more sophisticated interpolation scheme, more representative estimated displacements are generated in this case and the results improved, particularly the number of particles tracked (62% for the simple scheme versus 81% for the interpolative scheme). The discussion of the errors follows as above for the other subimage cases.

3.3.3 Divergent flows

To assess more accurately the errors incurred with each method for flows relevant to those studied as part of this thesis work (divergent free-surface flows), a simple numerical flow field was generated using a linear superposition of several stagnation-like flow elements and applied this flow to a pair of synthetic images. The image size in these tests was 672 px × 672 px, all other image characteristics remaining the same. The flow field was generated from the following velocity components:

$$\begin{aligned} u(x, y) &= \sum_{i=1}^n \frac{a_i}{2} \frac{x - x_{o_i}}{\sqrt{(x - x_{o_i})^2 + (y - y_{o_i})^2}} \\ v(x, y) &= \sum_{i=1}^n \frac{a_i}{2} \frac{y - y_{o_i}}{\sqrt{(x - x_{o_i})^2 + (y - y_{o_i})^2}}, \end{aligned} \quad (3.10)$$

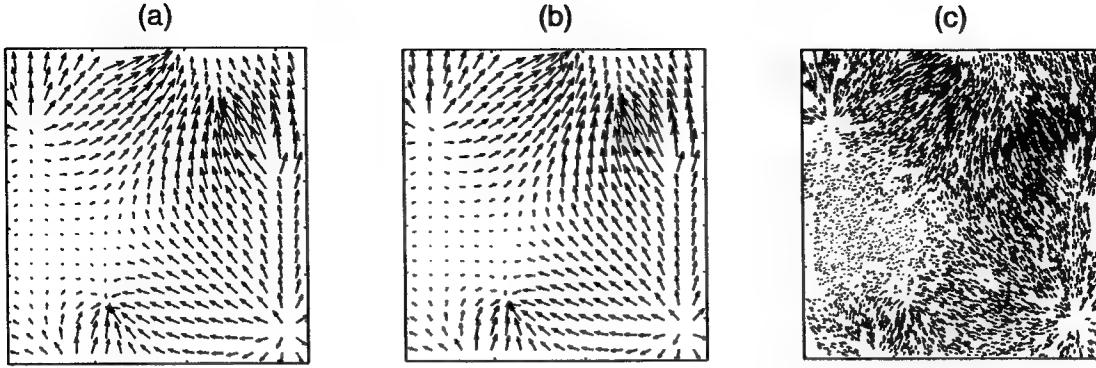


Figure 3-12: Displacement fields for divergent flow A: (a) prescribed displacement field, (b) hybrid DPIV result, and (c) raw PTV result.

where a_i is the i -th flow element strength, (x_{o_i}, y_{o_i}) is the i -th flow element center, and n is the number of elements used. The values for the flow element strengths and centers were generated randomly, and the strengths were allowed to be both positive (upwelling) or negative (downwelling). A value of $n = 8$ was used for each of the three flow situations generated. A sample prescribed flow is given in figure 3-12, and the velocity field results from the hybrid DPIV processing and the raw PTV processing are also shown. The error results for a 32×32 subimage size are shown in table 3.3. The RMS error and the mean relative error are defined as in equation 3.8 and equation 3.7 respectively, and the mean (relative) divergence error was defined as

$$E_{div} = \frac{1}{N} \sum_{i=1}^N \left| \frac{D_i - D_a}{\bar{D}} \right| \times 100\%. \quad (3.11)$$

The local “actual” divergence, D_a , was computed from a first-order finite difference of $(\partial u / \partial x) + (\partial v / \partial y)$, using the prescribed velocity field evaluated on the DPIV grid. The measured divergence, D_i , was calculated using the same finite differencing. The error was scaled by the mean magnitude of the divergence field,

$$\bar{D} = \frac{1}{N} \sum_{i=1}^N |D_a|. \quad (3.12)$$

(a) Flow A (maximum displacement = 6.1 px)

	FFT	DFFT	Spatial	Hybrid	PTV	PTV _{fine-grid}
E_{rms} (px)	0.131	0.115	0.107	0.107	0.087	0.110
E_{rel} (%)	3.77	2.36	1.84	1.85	1.61	2.09
E_{div} (%)	12.49	11.11	9.18	9.18	–	17.52

(b) Flow B (maximum displacement = 7.9 px)

	FFT	DFFT	Spatial	Hybrid	PTV	PTV _{fine-grid}
E_{rms} (px)	0.154	0.117	0.118	0.129	0.123	0.150
E_{rel} (%)	3.60	2.55	2.07	2.11	2.00	2.40
E_{div} (%)	11.00	8.17	7.49	7.52	–	17.30

(c) Flow C (maximum displacement = 9.8 px)

	FFT	DFFT	Spatial	Hybrid	PTV	PTV _{fine-grid}
E_{rms} (px)	0.189	0.097	0.086	0.086	0.163	0.171
E_{rel} (%)	2.83	1.59	1.44	1.44	1.41	1.80
E_{div} (%)	8.87	7.50	6.94	6.94	–	13.26

Table 3.3: Simulated divergent flow results. Boldface indicates best performance.

A mean value rather than the local value was chosen because in these flows the local divergence can become quite small, resulting in large relative errors that bias the mean relative divergence error. Particle tracking results are indicated as PTV, and reflect the performance of the raw, irregular data. The PTV_{fine-grid} results are the raw particle tracking results gridded to a fine resolution regular array as done for the rotational tests with an added post-processing step of passing the derived divergence through a 3×3 median filter. This filtering operation was performed in order to reduce the high wavenumber noise in the divergence fields as might be done in practice. As in the rotated image pair test, with 32×32 subimage size processing, the performance variability is not dramatic. Again, DPIV errors associated with the standard FFT scheme are the largest, and the spatial and hybrid methods exhibit the most accurate results in almost all cases. The PTV processing shows interesting performance. In terms of displacements, the raw PTV data show excellent accuracy, and the fine-grid PTV results, which are at a spatial resolution nearly 4 times that of the DPIV results, maintain the same level of accuracy. However, even in light of the

favorable accuracy in the displacement measurements, the higher resolution PTV results fall short in the calculation of divergence, showing relative errors twice that of the other methods. The increase in error is attributable to the same effect observed for the vorticity calculations in section 3.3.2. The reason that the increase is only a factor of two rather than the expected factor of four is a consequence of the median filter operation. In all cases, the greatest divergence errors are found to occur at the stagnation sources where the flow is very dynamic. Regardless of this increase in error, the higher resolution measurements provided by the $PTV_{fine-grid}$ are desirable when resolving small-scale flow features is important, or when measuring flows near boundaries or in regions of high shear.

3.4 Summary

These numerical experiments have shown that on a detailed level, the generation of the cross-correlation peak using the direct spatial correlation calculation provides highly accurate results from digital image data for a variety of flow fields. The method accomplishes this by incorporating more of the direct image content when computing the correlation surface, rather than relying on subimage extension and periodicity assumptions. The result is a more robust correlation peak. The dynamic FFT method is a viable processing scheme, but the standard FFT method has to be considered the least accurate of the group and its use would not be suggested unless speed is a primary concern. Both the DFFT and the hybrid methods can improve upon the accuracy of the FFT processing at only minimal computational cost. In fact, the hybrid method, which is an original approach to PIV data processing, was found to mimic the accuracy of the direct spatial method while only requiring a fraction of the computational effort, making it a useful contribution of this thesis work. The particle tracking routine studied has shown its ability to provide accurate displacement data at resolutions several times greater than standard DPIV techniques. Note that often DPIV data are oversampled (overlapping subimages), which increases the spatial frequency of measurement points. This does nothing, however, to reduce the size of the scales resolvable; the smallest scales measurable are still dictated by the size of the subimage.

These velocity extraction algorithms are an important element in the PIV system used and developed in this thesis. The image data from the forthcoming chapters were processed exclusively with either the direct spatial or the hybrid correlation methods for the final results. Particle tracking also was performed when the image data were well-suited for PTV analysis. On occasion, the FFT-based methods were used to efficiently verify the experimental arrangement, the image quality, and the particle displacements. In this manner, the five processing codes evaluated in this chapter comprise a useful suite of image velocimetry analysis algorithms for experimental fluid mechanics.

Chapter 4

A DPIV technique for free-surface flow measurements

4.1 Introduction

As shown in chapter 2, the critical physics of air–water mass transfer are confined to a very thin fluid layer near the air–water interface. With this layer being typically less than 1–2 mm in thickness, experimentally probing the flow structure in this region has proved exceedingly difficult. What often makes this task even more challenging is the propensity for a water surface to deform, most notably in the form of free-surface waves. When waves are present, the position of the air–water interface is a variable itself.

Several researchers have endeavored to tackle this problem. The efforts can be grouped loosely into two classes: near-surface profiling measurements, and near-/free-surface measurements parallel to the plane of the interface. The bulk of the efforts have focused on the first class of measurements and feature a variety of experimental tools. Two of the earliest studies belonging to this group were performed by McLeish and Putland [99] and Okuda *et al.* [105] in the 1970s. Both were based on the use of hydrogen bubbles to visualize and quantify the flow beneath surface water waves. Innovative as these studies were, Banner and Peirson [8] have pointed out several valid weaknesses of these approaches in accurately measuring the near-surface flow field. More recently, laser Doppler velocimetry (LDV) has been employed by Caulliez [16] and Melville and Rapp [100] to the free-surface wave problem.

Using a split-film anemometer probe, Brumley and Jirka [13] claim measurements of the velocity fluctuations to within $400 \mu\text{m}$ of the free surface for the case of a grid-stirred flow. The advent of the particle image velocimetry technique (and its derivatives) has spurred numerous other investigations into near-surface flow. These experiments have the advantage over some of the aforementioned ones in that they are non-invasive. Again, looking at wavy surfaces, Cowen *et al.* [21], Dieter *et al.* [29], Hering *et al.* [52], Hirsa *et al.* [55], Lin and Rockwell [84], and Peirson [106] have probed the very near-surface region with their PIV-based profiling methods, some even reporting measurements *at* the free surface. In these various techniques, the location of the free surface has sometimes been accounted for using accompanying wave-wire gauges and/or conditional sampling methods (for LDV) or by using image processing techniques (for PIV). In the latter, the interface between the air and water can appear as a thick (wavy) line. This presents difficulties in precisely determining the vertical position of the free surface, in addition to obscuring motions right at the interface (see figure 4-15a, for example). A novel approach to this dilemma was reported by Law *et al.* [81], in which the authors used fluorescent dye to accurately visualize the air–water boundary while simultaneously measuring the near-surface velocity field using PIV. The dye was selected such that when the measurement plane was illuminated with a light sheet, the water-side produced a luminance that contrasted sharply with the air-side. While the results of tracking the interfacial position appear favorable, the optical setup for this scheme is quite elaborate. Additionally, the degree of surface activity of the introduced dye is of obvious concern. In Law *et al.*, there is no mention of the properties of the dye.

Measurements near, or at, the free surface in planes parallel to the surface have been made using whole-field imaging-based methods. Measurements of the free-surface velocity field induced by the oblique interaction of a vortex ring with the surface have been reported by Weigand [138]. A similar approach was used by Logory *et al.* [89] in their work investigating the free-surface wake behind a vertical flat plate. Both of these studies report illuminating laser light sheets within 1 mm of the free surface, the thickness of the sheets also being approximately 1 mm. However, no mention is made in either study to explain how the effect of lensing by the fluid meniscus on the light sheet entering the water through the transparent tank walls was handled. This problem has been observed by the

present author. The meniscus is curved and any portion of the light sheet passing through it is refracted. Similarly, minor surface deformations can also distort the light sheet. A simple means to handle these problems has proved elusive. It is likely that in the two aforementioned studies, the light sheet is slightly below the interface.

Another approach to the whole-field imaging methods used at the free surface is illumination from above rather than in a plane close to the interface. Lloyd *et al.* [88] used diffused light from high-powered halogen lamps to illuminate floating particle tracers in a shallow-water flume flow. The particles, 5 mm in diameter, were painted black so as to contrast well with the white flume bed. Using a digital particle tracking technique, the surface flow around model islands was studied. While the large size of the particles in this study may have been acceptable (the imaged area was quite expansive—1.5 m × 1.3 m), such would not be the case for measurements of the fine-scale motions associated with interfacial transport. A similar technique was utilized by Tamburrino and Gulliver [127] in an open-channel flow. In this approach, photographic lamps were pulsed to briefly illuminate floating surface particles. The particles, which actually became small clusters due to their tendency to conglomerate, were imaged as streaks. The streak images were processed to yield individual particle velocities and the results interpolated to a regular grid for further processing. The most recent effort in this vein is that of Kumar *et al.* [77]. Using imagery of silver-coated micro-spheres illuminated by white light, this study put emphasis on free-surface features and qualitative descriptions, with a hierarchical DPIV algorithm used to make quantitative measurements as well.

In this chapter, a new technique for successfully probing the hydrodynamics of a free surface is presented. It is a whole-field DPIV imaging technique that is shown to measure the motion of the fluid within the upper 500 μm of the interface (a worst case estimate), in the plane of the free surface. A spatially intensive measurement in the plane of the interface was chosen over a profiling method for a number of reasons. First, as has been discussed above, locating the position of the free surface is a difficult task. Methods that attempt to measure the interfacial position have the potential pitfalls of an interface that is poorly defined and/or the image of the interface prohibits measurements very near the surface. Alternative approaches that can overcome such problems often will require sophisticated

optical arrangements, or as in the case of Law *et al.* [81], introduce possible surface contamination, which is, in general, not acceptable. Second, a profile measurement can only resolve the *surface* flow along a single dimension at most. In terms of the relationship between surface hydrodynamics and the processes of interfacial transport, a two-dimensional measurement of the free surface is more desirable, particularly when examining the role of surface divergence. In the proposed technique, laser light illumination of tracer particles was chosen rather than (diffuse) polychromatic light sources. This was done in order to (1) provide maximum particle illumination in the shortest amount of time so as to avoid streaking, and (2) allow for specialized optical filtering of the imaging camera. This latter reason is explained in section 4.2; a notable benefit of such filtering is that it enables this technique to work successfully for low-amplitude, undulating water surfaces (see chapter 5).

4.2 Description of the surface PIV technique

The free-surface flow measurement technique developed as part of this thesis work is relatively straightforward in concept and in execution. Its uniqueness relies on three components: (1) monochromatic laser light, (2) fluorescent PIV seed particles, and (3) optical filtering of the imaging camera. The hardware components are those of any typical DPIV system. Figure 4-1 illustrates the fundamental elements of the setup. The aspects of the system that are specifically relevant to the surface measurement technique are detailed in the following few sections; discussion of the other components can be found in section 4.4.3.

4.2.1 Laser illumination

The laser used in this thesis work was a commercially available, dual, frequency-doubled Nd:YAG laser system. Two pulsed Nd:YAG infrared laser heads are mounted on a single baseplate. The 1064 nm beams are polarization-combined and then enter a second harmonic generator to produce visible green light at 532 nm. Dichroic mirrors separate the visible from the residual infrared light and direct the beam to the experiment. The light emerging from the laser is vertically polarized. Each laser cavity is capable of emitting 3–5 ns pulses

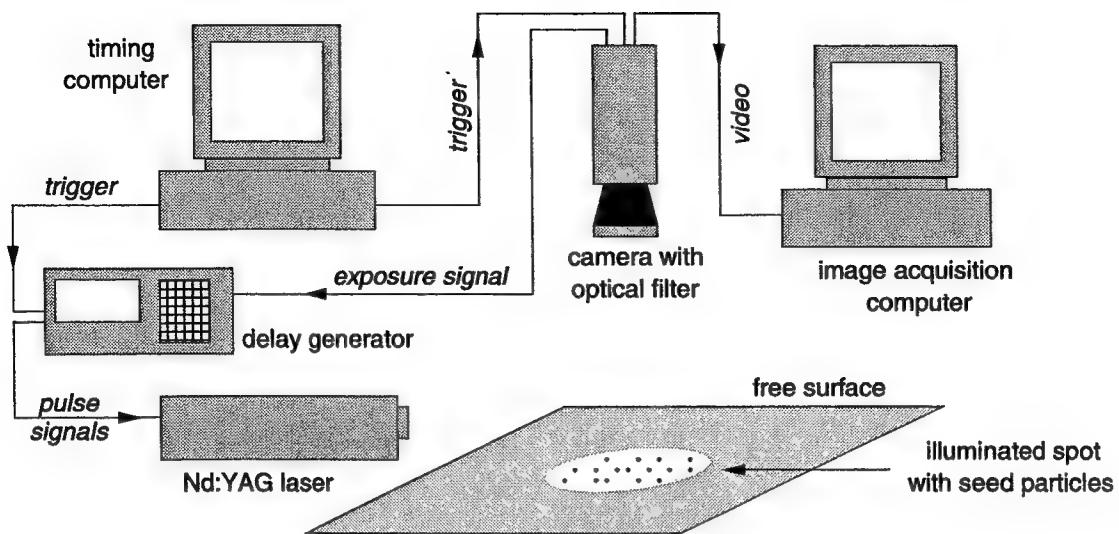


Figure 4-1: Basic elements of the (surface) PIV system developed in this thesis.

at a maximum repetition rate of 15 Hz. The energy per pulse typically was 100–120 mJ. For comparison, a 5 W continuous laser would require about 22 seconds to deliver an equal amount of light energy; this length of exposure time would result in a significant amount of particle image streaking for most flows. The main head unit is attached to two individual power supply/cooling units via umbilicals. This allows for easy positioning of the beam source. Each laser is externally triggerable and can be synchronized with the system camera.

4.2.2 Seed particles

All surface measurement flows were seeded with acrylic spheres impregnated with Rhodamine 6G and dichloro-fluorescein dyes. When these particles are excited with light in the range 511–532 nm, they emit in the range 570–595 nm depending on the excitation. Using a spectrophotometer, the emittance wavelength based on a Nd:YLF laser excitation (523.4 nm) was estimated to be roughly 565 nm. The particles used for this thesis were either 20–40 μm or 80–120 μm in diameter, and had a specific gravity of 0.92–1.1.

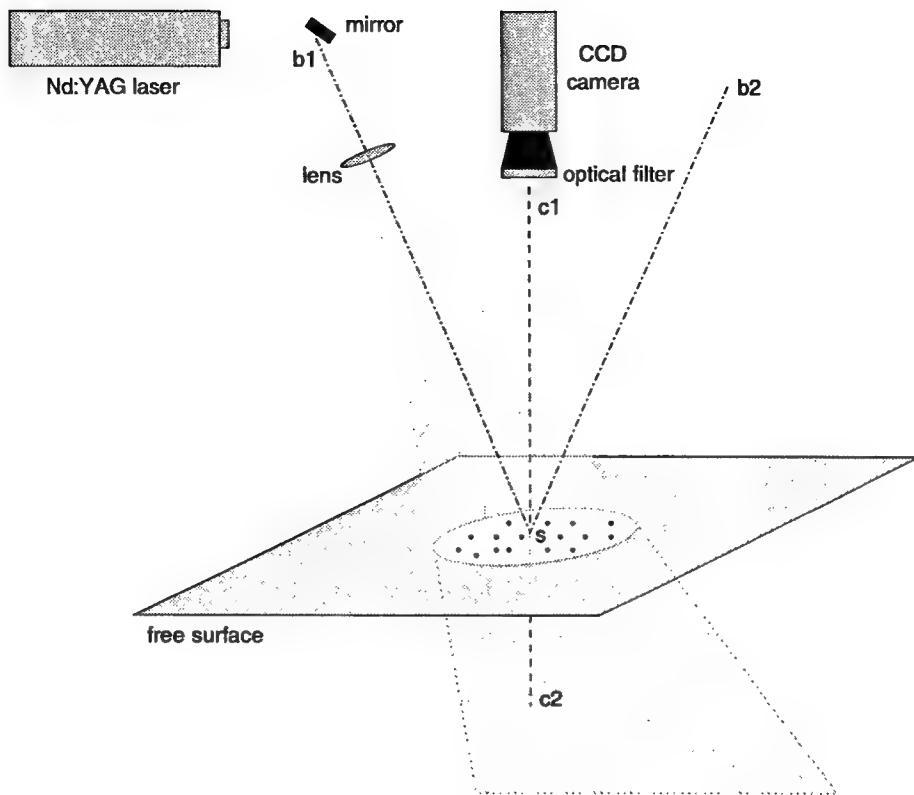


Figure 4-2: PIV surface flow measurement optical arrangement. The laser beam is redirected by a high energy mirror along beam axis b1-s. A spherical lens fans the beam into a three-dimensional, diverging light cone incident on the free surface. The camera optical axis is c1-s-c2. This axis lies in the plane of incidence formed by b1-s-b2.

4.2.3 Optics

Figure 4-2 shows the optical arrangement used in the surface measurement technique. A high energy solid-state laser mirror was used to direct the beams from the laser downward, at an angle, toward the free surface. After this redirection, the beams passed through a fused silica, plano-concave spherical lens that transformed the 5 mm diameter collimated beam into a diverging cone of light incident on the free surface. From the specified focal distance of the lens (56 mm for this wavelength) the diverging angle of the cone was computed to be 2.56° . For the arrangement typically used, the lens was positioned such that a light spot of nominally 8 cm in diameter was formed at the free surface. The camera (see section 4.4.3) was located in the plane of incidence, formed by the incident and reflected

light cone axes. When fluorescent particles are at the water surface, the light reflected from the interface is dominated by two wavelengths: the laser light at 532 nm, and the particle fluorescence emission presumably near 565 nm. The reflected laser light, while being only a small fraction of the incident light, is so intense it completely saturates the camera sensor. (When at nominal laser power, the reflected light permanently damages the image sensor!) To remove this unwanted reflection, an interference bandpass filter was used. The filter was selected so as to pass only the wavelength of the particle fluorescence emission. Two precision optical filters were considered, each having a passband width of 10 ± 2 nm. One filter was centered at 560 nm, the other at 570 nm. Each is specified by the manufacturer to have a minimum energy transmittance (at the passband) of 50%; however, values significantly higher reportedly are possible. By bandpass filtering the camera, the laser light is eliminated from the video image and only the particle images are detected.

A potential concern with this approach is the fact that the cone of laser light penetrates the water column, illuminating particles at depth as well as at the surface. The lens employed with the camera had an estimated depth of field of 0.5–1.0 cm for the *f*-number(s) used. Thus, particles in the upper centimeter or so of the water can be in focus; those below this depth will be increasingly de-focused, resulting in broader, dimmer particle images. However, it was found that two factors contributed to eliminate this negative consequence of the illumination scheme. First, it was observed that particles at the very surface of the water appeared much brighter to the camera sensor than did particles at depth. Second, because of this intensity discrepancy, aside from blocking the unwanted reflected laser light, the bandpass filter also helped to optically threshold the particle images. This was due to the filter having a finite amount of transmission loss (discussed at the conclusion of section 4.3.3), and the camera CCD sensor having a finite sensitivity. Bright particles (at the surface) had filtered intensities above the camera minimum sensitivity, while dimmer particles (at depth) were below the camera sensitivity limit. Consequently, these dimmer particles below the surface avoided detection by the camera sensor. The following section provides some theoretical basis for this observed phenomena and additional physical insight. More critically, sections 4.4–4.6 describe the experimental validation of the technique.

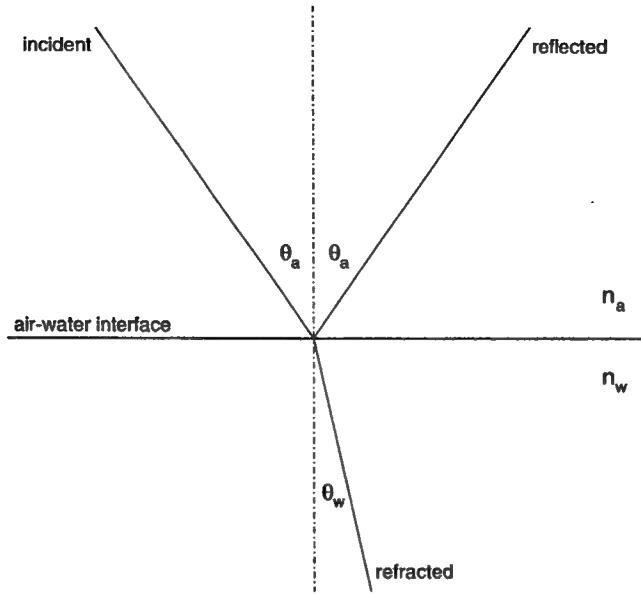


Figure 4-3: Reflection and refraction of light at an air–water interface.

4.3 Theoretical context

To explain the discrepancy in brightness of particles at the surface and those below, a number of possible explanations were explored. They were: (1) the effect of reflected light at the free surface, (2) attenuation of light with depth in the water, and (3) scattering behavior of particles in air versus that in water. The basis for each explanation lies with the assumption that particles at the surface actually are floating on the surface such that some fraction of their surface area is exposed to the air.

4.3.1 Reflection at an air–water interface

Particles on the water surface can be considered to receive all of the incoming laser light. Those below the surface will receive the amount of light that is transmitted through the air–water interface after a finite amount of light is lost to surface reflection. Considering figure 4-3, the amount of incident light that is reflected by a water surface is calculated

using the Fresnel reflectance formulae (Hecht [51]):

$$R_{\perp} = \frac{\sin^2(\theta_a - \theta_w)}{\sin^2(\theta_a + \theta_w)}, \quad (4.1)$$

$$R_{\parallel} = \frac{\tan^2(\theta_a - \theta_w)}{\tan^2(\theta_a + \theta_w)}. \quad (4.2)$$

In these expressions, θ_a is the angle of incidence (in air) and θ_w is the angle of refraction (in water). These are related through Snell's Law,

$$n_a \sin \theta_a = n_w \sin \theta_w, \quad (4.3)$$

where n_a and n_w are the indices of refraction for air and water, respectively. The reflectance coefficient R_{\perp} represents the reflected fraction of incident light that is polarized perpendicular to the incident plane formed by the incident and reflected rays, and R_{\parallel} represents the reflected fraction of incident light that is polarized parallel to the incident plane. In the optical arrangement depicted in figure 4-2, the incident light is oriented with parallel polarization, thus equation 4.2 can be used to estimate the amount of incident laser light lost to reflection at the surface. For $\theta_a \approx 15^\circ\text{--}20^\circ$, which corresponds to the actual physical range used, and taking $n_a = 1.000$ and $n_w = 1.335$ for 532 nm light (Shifrin [121]), only 2.3–2.5% of the incoming laser light is lost to reflection. This would imply that particles below the free surface effectively receive the same amount of incident light energy. Hence, surface reflection cannot explain the particle image intensity discrepancy.

4.3.2 Light attenuation in water

The attenuation of light in water can be described by a (total) attenuation coefficient that captures the combined effects of absorption and scattering in the water, the former effect dominating. For pure water, the attenuation coefficient for light near 532 nm has been found to be approximately 0.05 m^{-1} (Jerlov [65]). At a depth of 10 cm, light at this wavelength is attenuated by only 0.5%. Therefore, attenuation effects also can be ruled out as an explanation.

4.3.3 Scattering of light by particles

Both Adrian [3] and Raffel *et al.* [108] have pointed out that the ratio of the refractive index of a particle to that of its surrounding medium is a significant parameter in the amount of light scattered by the particle. Because the refractive index of water is 1.33 times that of air, particles in air are found to scatter considerably more light than particles in water. To explore this behavior for the surface PIV setup, the general scattering theory of Mie (1908) was employed. (See Kerker [69] for details of Mie's theory.)

Mie scattering theory addresses the exact solution of the scattering of a plane electromagnetic wave by an isotropic, homogeneous sphere of arbitrary size. From the far-field solution of the Mie formulation, the intensity of the energy scattered by a spherical particle of radius a and refractive index m_1 in a medium having refractive index m_2 is determined by the amplitude functions for perpendicular and parallel polarization light,

$$S_1(\theta) = \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} \{a_n \pi_n(\cos \theta) + b_n \tau_n(\cos \theta)\} \quad (4.4)$$

$$S_2(\theta) = \sum_{n=1}^{\infty} \frac{2n+1}{n(n+1)} \{a_n \tau_n(\cos \theta) + b_n \pi_n(\cos \theta)\}, \quad (4.5)$$

respectively, where θ is the scattering angle, measured from the direction of the incident wave propagation. The angular functions in equations 4.4 and 4.5 are given by

$$\pi_n(\cos \theta) = \frac{P_n^{(1)}(\cos \theta)}{\sin \theta} \quad (4.6)$$

$$\tau_n(\cos \theta) = \frac{d}{d\theta} P_n^{(1)}(\cos \theta), \quad (4.7)$$

where $P_n^{(1)}(\cos \theta)$ are the associated Legendre functions of degree n and order 1. The (complex) coefficients a_n and b_n are given by

$$a_n = \frac{\psi_n(\alpha)\psi'_n(\beta) - m\psi_n(\beta)\psi'_n(\alpha)}{\zeta_n(\alpha)\psi'_n(\beta) - m\psi_n(\beta)\zeta'_n(\alpha)} \quad (4.8)$$

$$b_n = \frac{m\psi_n(\alpha)\psi'_n(\beta) - \psi_n(\beta)\psi'_n(\alpha)}{m\zeta_n(\alpha)\psi'_n(\beta) - \psi_n(\beta)\zeta'_n(\alpha)}, \quad (4.9)$$

where the following dimensionless parameters are prescribed

$$\alpha = 2\pi a/\lambda = 2\pi m_2 a/\lambda_0 \quad (4.10)$$

$$\beta = m\alpha, \quad m = m_1/m_2. \quad (4.11)$$

The wavelength λ_0 is that of the incident wave in a vacuum. Appearing in equations 4.8 and 4.9 are the Riccati-Bessel functions, defined as

$$\psi_n(x) = \left(\frac{\pi}{2}x\right)^{1/2} J_{n+\frac{1}{2}}(x) \quad (4.12)$$

$$\chi_n(x) = -\left(\frac{\pi}{2}x\right)^{1/2} N_{n+\frac{1}{2}}(x), \quad (4.13)$$

where $J_{n+\frac{1}{2}}(x)$ and $N_{n+\frac{1}{2}}(x)$ are the half integral order Bessel and Neumann functions, respectively. Also needed is the linear combination,

$$\zeta_n(x) = \psi_n(x) + i\chi_n(x). \quad (4.14)$$

Primed quantities in equations 4.8 and 4.9 denote differentiation with respect to their arguments. In terms of simplifying computations, it is useful to invoke one of the standard recurrence relations for all cylinder functions,

$$\frac{d}{dx} Z_n(x) = \frac{1}{2} Z_{n-1}(x) - \frac{1}{2} Z_{n+1}(x). \quad (4.15)$$

For an incident beam of unit intensity polarized perpendicular to the scattering plane (formed by the incident direction and the direction of the scattered wave), the intensity of scattered light is given by

$$I_1(\theta) = \frac{\lambda^2}{4\pi^2 r^2} |S_1(\theta)|^2 = \frac{\lambda^2}{4\pi^2 r^2} i_1. \quad (4.16)$$

For a beam polarized parallel to the scattering plane, the scattered intensity is

$$I_2(\theta) = \frac{\lambda^2}{4\pi^2 r^2} |S_2(\theta)|^2 = \frac{\lambda^2}{4\pi^2 r^2} i_2. \quad (4.17)$$

Parameter	Value
laser wavelength, λ_0	532 nm
refractive index, m_1	1.500
refractive index, m_2 (air)	1.000
refractive index, m_2 (water)	1.335
particle diameter, $2a$	1–30 μm

Table 4.1: Parameters used in theoretical Mie scattering calculations.

In these last two equations, r is the radial distance from the particle center.

To evaluate the scattering behavior of particles on the water surface (in air) versus particles below the interface (in water), a numerical implementation of the above mathematical formulation was developed. Table 4.1 summarizes the relevant parameters used in the calculations. The infinite series in the expressions for S_1 and S_2 have been found to converge for $n \approx 1.2\alpha$ (Shifrin [121]). Based on this criterion, and given the limitation of the numerical code (Legendre functions only up to degree 255 could be easily computed), an upper limit on the particle diameter able to be studied emerged: 30 μm . Since the size of the particles actually used included cases of 20–40 μm , this limit is not overly prohibitive.

A Mie scattering diagram for acrylic particles of diameter 30 μm is shown in figure 4-4. Only i_2 is shown since the two intensity functions are very similar for this case. Furthermore, i_2 represents the scattering that arises from a beam polarized parallel to the scattering plane, and in the arrangement of figure 4-2, this is the case for the illuminating laser light. Therefore, i_2 should be the dominant contribution to the scattered intensity observed by the camera. The dependence of the scattered intensity on particle diameter is shown in figure 4-5. As can be seen from these two figures, the intensity of light scattered by a particle in air is notably greater than the scattering for the same particle in water. This would seem to support the observed phenomena of surface particles appearing much brighter than those in the bulk. However, the present Mie calculations are simplified in that they do not take into account the fluorescence of the PIV particles or the fact that the particles are not pure acrylic.

Since the difference in intensity between particles in air and those in water was found to be approximately a factor of 7, the combination of the finite transmission loss of the optical

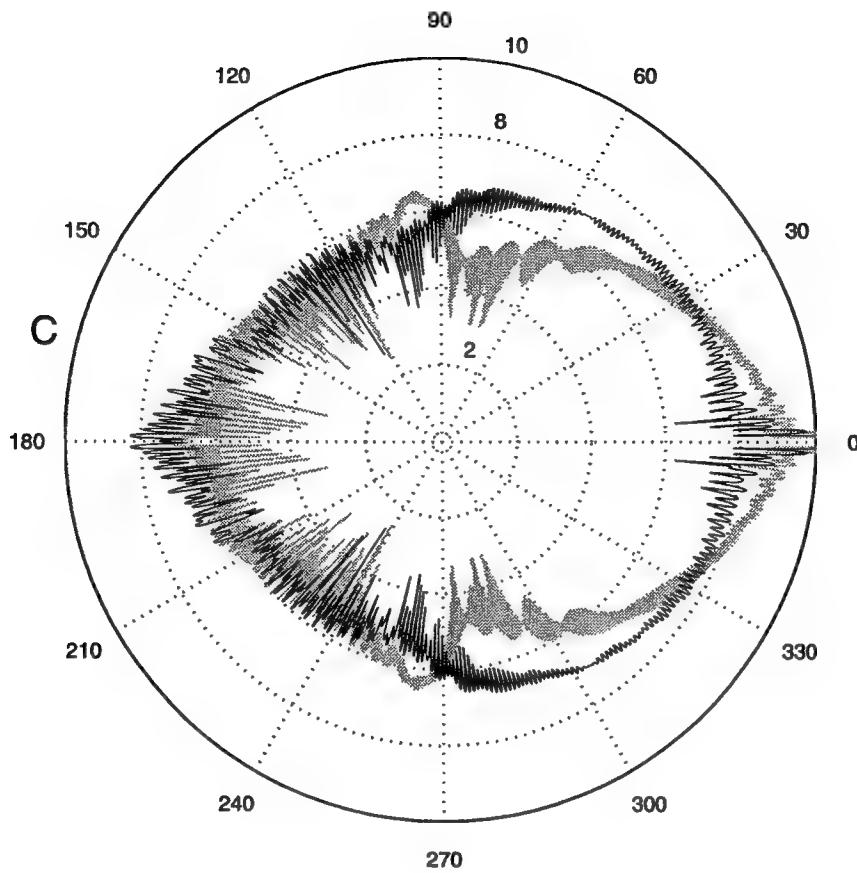


Figure 4-4: Mie scattering diagram for a particle with a diameter of $30 \mu\text{m}$ and a refractive index of 1.50. Plotted is the relative intensity i_2 for a particle in air (black line) and a particle in water (gray line). The intensity scale is plotted logarithmically with adjacent scales differing by a factor of 10^2 . The incident beam is traveling from $\theta = 180^\circ$ toward $\theta = 0^\circ$. The relative location of the imaging camera is denoted by "C."

filter and the finite sensitivity of the camera sensor could provide optical thresholding of the particle images. To test this, a small tank was set up with the PIV seed particles in water. The fluid and particles were illuminated in a narrow plane with a laser light sheet using the Nd:YAG system. Video images were collected using the PIV system camera with and without the two optical filters in place, keeping all other conditions the same. The attenuation due to the filters was dramatic. The mean intensity of the image taken without optical filtering was 188.9. The mean intensity of the image obtained using the 560 nm filter was 12.7, and the mean intensity for the 570 nm filter was 9.5. The difference between

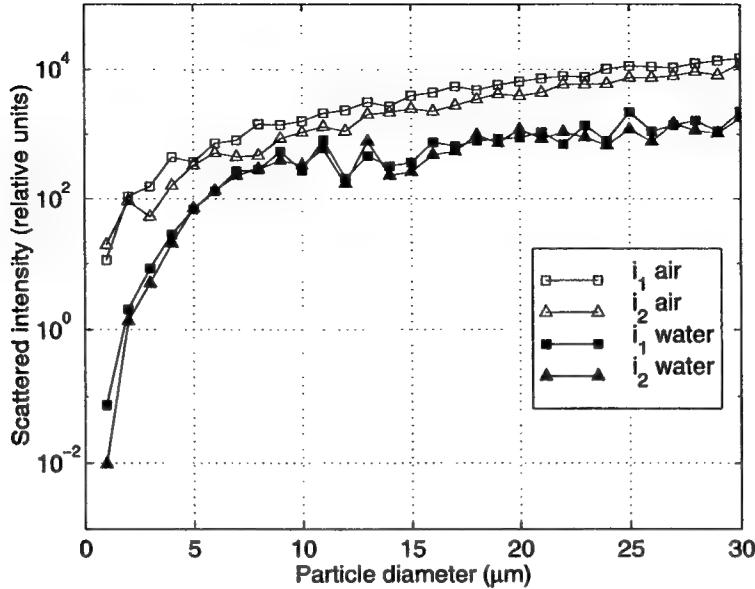


Figure 4-5: Dependence of scattered intensity on particle diameter. Shown are the relative scattered intensities i_1 and i_2 averaged over the scattering angle range of 150° – 180° .

a filtered camera image and an unfiltered camera image is a factor of 20! This implied that the PIV particles fluoresce at a wavelength other than the estimated 565 nm, and the particle emission is attenuated by a significant amount. This result works in favor of the surface PIV measurement technique—the light intensity emitted from the PIV particles in the bulk is too weak to be detected by the bandpassed camera optics. Only *high intensity* light is sensed by the camera CCD after passing through the filter. Therefore, only surface particles, which are potentially an order of magnitude more intense, are able to be detected by the filtered camera.

A final test of the influence of subsurface particles was performed with the numerically generated images from chapter 3. The image pair that produced a 5.25 px uniform translation of particles (denoted $\mathcal{A}_1, \mathcal{B}_1$) was modified to simulate a sheared near-surface flow. The 5.25 px pair was selected as a characteristic pixel displacement since in the majority of the PIV work in this thesis, a 5 px displacement was a representative value. The modification involved taking another randomly seeded image with particles as specified in section 3.3, and reducing all particle intensities uniformly by a factor of 7 (the result denoted \mathcal{A}_2). From

this image, a displaced image (denoted \mathcal{B}_2) was created that yielded a -1.0 px uniform translation of particles. A final image pair was formed as $\mathcal{A} = \mathcal{A}_1 + \mathcal{A}_2$ and $\mathcal{B} = \mathcal{B}_1 + \mathcal{B}_2$. This image pair, therefore, might be considered a simulation of the case where the flow at the very surface was moving in one direction with velocity V , and the flow below the surface was moving in the opposite direction with velocity $0.19V$. The pair \mathcal{A}, \mathcal{B} was processed with the hybrid correlation exactly as was done for the chapter 3 results. From chapter 3, the mean displacement found using the original images $\mathcal{A}_1, \mathcal{B}_1$ was 5.2508 px, with an RMS error of 0.0092 px and a mean bias error of 0.0007 px. The result from processing \mathcal{A}, \mathcal{B} was a mean displacement of 5.2496 px, an RMS error of 0.0185 px, and a mean bias error of -0.0004 px. The effect of the dimmer particles “flowing” in the opposite direction is only notably apparent in the RMS errors. A factor of 2 increase in the RMS error is found for the simulated shear case over the purely uniform flow. The differences in the means for each case are exceedingly small.

The results from the theoretical Mie scattering analysis, the optical filtering experiment, and the flow simulation test just described all point to a surface PIV velocity measurement that should be nearly free of the effects of particles at depth. However, each of these results was based on numerous assumptions and are only possible indicators of reality. To verify accurately the surface PIV technique, actual experiments needed to be performed that could genuinely assess the validity of this proposed method.

4.4 Validation of the technique: Phase 1

Validating this free-surface flow measurement technique is a non-trivial task. The major issue to be resolved is to what degree, if any, do particles at depth contribute to the measurement. Therefore, a test flow that exhibits shear at the surface is desirable since the velocity beneath the surface will be different from that at the surface. If the technique is indeed integrating the flow over a finite depth, this will be revealed in a discrepancy between the flow surface velocity and the measured velocity. Generating such a flow is not difficult; however, *accurately* knowing the surface velocity of the prescribed flow is quite difficult.

In this and the following two sections, the technique outlined in section 4.2 is validated

through a sequence of comparative experiments in a paddle-driven racetrack flume. In the first phase of the validation, results are presented from a groundtruthing experiment that compared quasi-coincident bulk flow measurements from an independent velocimetry system and the DPIV system used in this thesis work. The second phase entails a similar comparison, but the measurements for this experiment were made at the free surface, with an unsheared velocity profile at the surface. In the third phase, a second free-surface comparison was made. In this case, the flow was wind-driven, creating a sheared velocity profile at the interface, and providing the true test of the surface PIV technique. For each comparison, attention to precise accuracy is given, as small errors can have significant consequences for this free-surface measurement technique.

The first step in validating the free-surface flow measurement technique was to ground-truth the DPIV system operating in its standard, bulk flow measurement configuration. This enabled meaningful interpretations of the PIV results from the subsequent experiments. The “system” includes both the hardware implementation and the image processing component described in chapter 3. To ascertain the performance of the DPIV system as a whole, comparative measurements were made with a laser Doppler velocimeter (LDV).

4.4.1 Experimental facility

All of the technique validation experiments were performed in a paddle-driven racetrack flume located in the Rinehart Coastal Research Center at the Woods Hole Oceanographic Institution. The water is circulated by a series of vertical paddles that cycle through one of the two linear segments of the flume. The flume has a linear raceway section 7.6 m long, 76 cm wide, and 25 cm deep. The side walls are made of glass for complete optical access. The bottom is PVC and has a 100 cm by 50 cm glass plate at the test section, which is located near the end of main raceway. Flow speeds in the range 2 cm/s to 55 cm/s are achievable. For all experiments, filtered sea water was used. The ambient particles in this water served adequately as LDV scattering material (see below) and as flow tracers for bulk flow PIV.

4.4.2 Laser Doppler velocimeter

The LDV used for the comparison experiments was a two-axis, backscatter device. In a backscatter LDV, all of the transmission optics and receiving optics are integrated into a single head unit. This is an advantage over traditional forward scatter LDVs that require separate units and must be properly aligned. The LDV measures flow velocities along two axes using two split-beam pairs that intersect at the measurement volume. The sample volume is ellipsoidal, with a long axis in the beam direction of 4.19 mm in length and a circular cross-section of 197 μm in diameter. The system works by dividing a laser beam in two and forcing the resulting two coherent beams to cross in the flow. Microscopic fringe patterns are formed in the intersection volume. Small seed particles in the flow pass through the fringe pattern and scatter light back toward the detector optics. The scattered light contains a Doppler shift—the Doppler frequency—which is proportional to the velocity component perpendicular to the bisector of the two laser beams. With a known wavelength of the laser light and a known angle between the intersecting beams, a conversion factor between the Doppler frequency and the velocity can be calculated. Additional components of velocity can be measured using one or two more beam pairs of different wavelengths.

4.4.3 Digital particle image velocimeter

The hardware for the DPIV system consists of a laser, seed particles, a digital video camera, a computer to acquire and store the video images, a delay generator, and a computer to control the laser and camera timing schemes when necessary. The laser and seed particles have been discussed in section 4.2. To image the flow field, a high resolution, digital CCD (charge coupled device) camera was used. The camera features an interline transfer sensor that has a pixel array of 1008 (C) \times 1018 (R) with progressive scan readout yielding full-frame output (both fields captured simultaneously). Each pixel element of the CCD array is 9 $\mu\text{m} \times 9 \mu\text{m}$. The camera can operate in several different exposure modes, of which, two are relevant here. In the first, referred to as “continuous” mode, the camera runs at 29.009 Hz, exposing each frame for nominally 33 ms. This was used as the standard mode for capturing continuous sequences of images. In the second mode, “double exposure” mode, the camera takes one image, exposed for 1–255 μs , transfers it to an on-board buffer, and

then 1–5 μ s later, acquires a second image, which is exposed for 33 ms. The length of the first exposure and the delay between exposures is selectable by the user. In this dual exposure mode, the camera can be left to time itself, running at the nominal 30 frames/sec (15 frame pairs/sec), or it can be externally triggered up to the 30 frames/sec limit. The camera outputs digital, 8-bit images.

To acquire the camera video, a computer equipped with a digital video capture board was used. The board serves to collect the image data as it comes in and then streams it to system RAM. The imaging computer was outfitted with 256 Mb RAM; this was later increased to 768 Mb. Accompanying software interfaced with the camera and allowed for image sequence capture, basic analysis, and image storage. Stored images were written ultimately to CD-ROM media.

A delay generator was used to properly time the two laser pulses with respect to the camera video signal. When the delay generator was triggered, by either the camera start-of-exposure signal or the system timing computer, it was set to delay a specified time, then trigger the first laser, delay another specified amount of time (into the next video frame), and then trigger the second laser. The time between laser triggers becomes the time between flow field exposures, denoted Δt .

For certain situations, a second computer also was part of the PIV system. This computer performed camera and laser timing synchronization. The computer was equipped with a I/O board that was programmed to generate specific timing schemes. This allowed for laser pulse generation and camera exposure synchronization operations.

4.4.4 Experimental arrangement and procedure

For the groundtruthing experiment, the two flow measurement systems were arranged as shown in figure 4-6. The region of interest was located roughly 6 m downstream from the start of the linear raceway section. The flow depth was 18 cm and the nominal flow speed was 15 cm/s. Ambient particles in the sea water were used as both LDV scatterers and PIV flow tracers. Both systems were set up such that profile measurements of the streamwise velocity could be made at several different heights above the bottom. The LDV head unit was mounted to a vertical translation stage thus enabling profiling; the PIV camera position

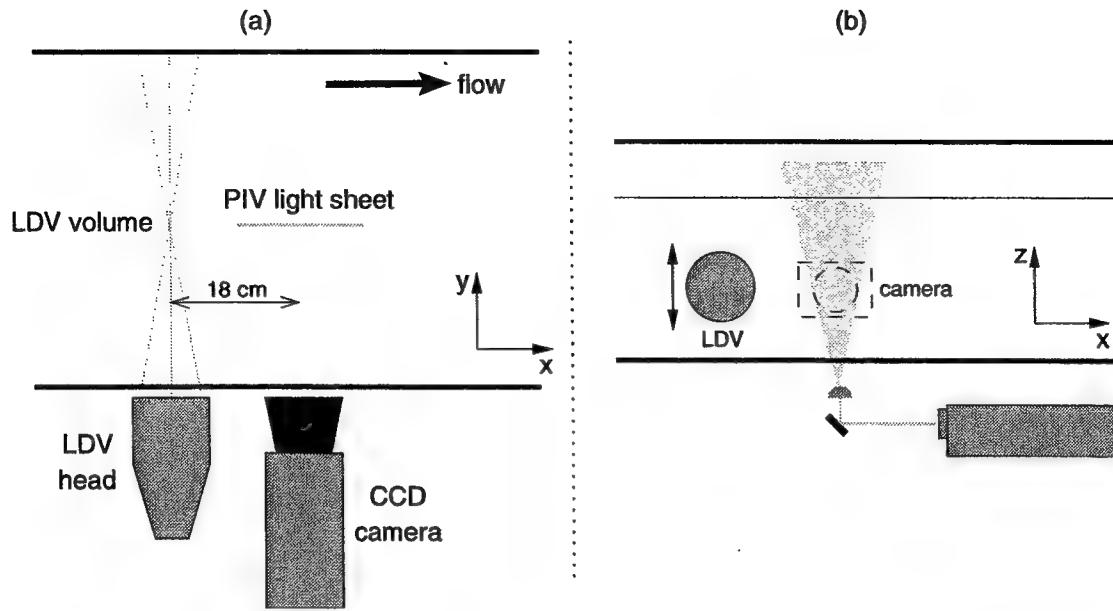


Figure 4-6: Experimental setup for the phase 1 comparison between LDV and standard PIV systems studying a channel flow. (a) Plan view of the test region. (b) Side view of the same region. Not to scale.

was fixed. The LDV optical axis was normal to the glass side wall and the measurement volume was located at mid-channel. Both measurement axes were enabled. The PIV system was positioned just downstream of the LDV, as near as physically allowable. The Nd:YAG laser was positioned below the channel and a light sheet was formed using a Powell lens. The sheet was sent vertically through the glass bottom in a wall-parallel plane at mid-channel. The camera was oriented normal to the side wall and focused on the illuminated plane. The image acquisition was set up such that images of dimensions 172×1018 were acquired in continuous mode. The corresponding physical domain was 9.7 cm in the vertical direction by 1.6 cm in the streamwise direction. The lower boundary of the image corresponded to 5 mm above the flume bottom. The streamwise distance between the LDV and the PIV measurements was roughly 18 cm. The time separation between laser pulses was 4 ms; streamwise particle image displacements were on the order of 6 px.

Measurements using both systems were made simultaneously. Synchronization of the two measurements was handled during the data processing. The duration of each measure-

ment was limited to the maximum amount of video that could be acquired with the PIV system. This storage limit translated into 1128 images, or roughly 40 seconds of PIV data. One minute records were taken with the LDV, with LDV data rates of 48 Hz on average. A coordinate system was chosen such that x measured the streamwise direction, and z measured the vertical direction with positive upward and zero at the flume bed. Measurements were collected at $z = 2, 3, 4, \dots, 9$ cm. While the PIV setup enabled a measurement of the complete profile instantaneously, only the velocity at the height corresponding to the LDV position was used for comparison in the analysis.

The experiment described above served as the main focus of the groundtruthoring effort. Additionally, however, a measurement was made to compare velocity spectra results for the two instruments. The limited PIV record length in the above setup prevented a good comparison of spectral properties. To obtain PIV spectra, the video acquisition system was reconfigured to capture only a 96×96 section of the camera image. This allowed for 10 minutes of PIV velocity data to be stored. For this measurement, data was successfully collected for both the LDV and PIV at a height above bottom of 2 cm. The experimental arrangement was identical to that already described. In this case, however, the two devices were synchronized to begin acquisition to within 300 ms of each other. The PIV processing was done so as to yield a single two-component velocity vector from each image pair. The spectral results are presented at the conclusion of the next section.

4.4.5 Analysis and results

Each pair of PIV-LDV time series measurements represented a single comparison at a particular z location. The processing of each time series pair was identical for all cases. The PIV data were processed using the hybrid correlation method discussed in chapter 3. Subimages of dimensions $32 \text{ px} \times 32 \text{ px}$ were used. This yielded a measurement volume that was $3 \text{ mm} \times 3 \text{ mm}$ in the xz -plane and $1\text{--}2 \text{ mm}$ in the cross-stream direction, the latter based on the laser sheet thickness. Velocity vectors were computed every 8 px (0.8 mm) in the vertical and every 16 px (1.5 mm) in the horizontal. Outlier removal was enabled, but no further post-processing was performed. A two-component velocity time series at 14.5 Hz was formed by sampling the PIV field at a single streamwise location at the appropriate

height above the bottom. In a few cases, this involved interpolating between two vertical locations. In this thesis, the following decomposition is used for the velocity:

$$\tilde{u}(t) = u + U \quad (4.18)$$

$$\tilde{v}(t) = v + V \quad (4.19)$$

$$\tilde{w}(t) = w + W \quad (4.20)$$

where $\tilde{u}(t), \tilde{v}(t), \tilde{w}(t)$ represent the three components of the fluid velocity, u, v, w represent the (turbulent) fluctuating components of the velocity, and U, V, W , are the respective means of each of the velocities. Mean values may be either temporally averaged, spatially averaged, or ensemble averaged, and are noted as such in the text.

Since the two devices measured at different streamwise positions and the sampling period was rather short, it was necessary to locate the segments of the time series that reflected measurements of the “same” flow. This was done by correlating the two series to determine a temporal lag between the two data sets. This is illustrated in figure 4-7. Once the “coincident” segments were determined, the data were corrected for possible angular misalignment. Rotation in the xz -plane was done such that the mean vertical velocity was equal to zero. The misalignments were computed to be 0.8° and 0.3° for the PIV and LDV, respectively. The corrections based on these rotations were insignificant.

When initially comparing the mean velocities (horizontal and vertical) from the two systems, the agreement was found to be excellent. However, a systematic bias between the two mean horizontal velocity estimates was discovered; the PIV data consistently underestimated the horizontal velocity component as compared to the LDV. This mean bias was found to be 0.14 cm/s, or 0.9% relative to the LDV measurement. A mean bias was not detected for the vertical components of velocity. Possible PIV sources of this discrepancy include spatial averaging of the sample volume, cross-correlation mean bias error, physical scale calibration error, and measurement height mismatch. The bias due to the spatial averaging of the PIV processing would incur errors less than 0.02% based on a logarithmic velocity profile that was fit to a measured PIV profile. This is well below the observed error. Considering the correlation calculation, the bias of 0.14 cm/s maps to a displacement bias

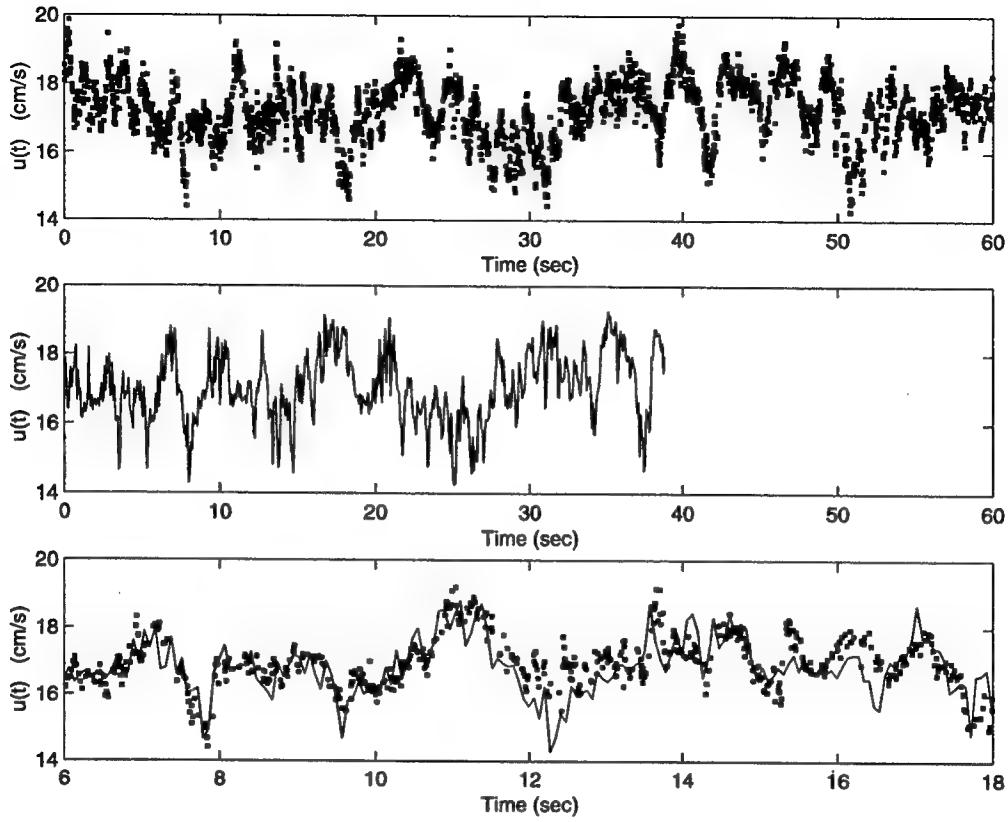


Figure 4-7: Horizontal velocity time series $\tilde{u}(t)$ showing location of “coincident” samples. (a) LDV complete time series, (b) PIV complete time series, and (c) LDV (points) and PIV (line) coincident data, shown over abbreviated time scale for clarity. The PIV data lag the LDV data by 4.3 seconds. Note, the two time series do not begin at the same time.

of 0.06 px. The performance analysis of chapter 3 indicated a possible bias magnitude of 0.0006 px for the hybrid correlation scheme. While the latter result derives from ideal, synthetic images, it is unlikely that actual images would result in an amplification of this bias error by two orders of magnitude. To calibrate the PIV images to physical units, a number of images of a scale are collected for analysis. Using the scale images, a known physical length is measured in pixel units, and a calibration factor determined. Values from all scale images are averaged to arrive at the final calibration factor used in the processing. Given the pixel displacements for this flow, an error in measuring the known length of about 8 px would be required to account for the bias observed. It is believed that errors of at most 2 px are possible. This would still only account for 25% of the discrepancy. Determining

the proper vertical locations in the PIV images was accomplished by interrogating images of a scale placed vertically in the measurement plane. The image locations of the various physical elevations are believed to be accurate to within 1/2 px. The sampled PIV data are located within ± 1.5 px of these locations. In physical units, a height mismatch of 2 px is only 0.19 mm. The largest variation of the flow velocity over such a small distance was estimated as 0.016 cm/s, only 11% of the observed bias.

As a perfect velocity measurement does not exist, it is of course possible that the LDV overestimated the flow velocity. Possible evidence for such was reported by Voulgaris and Trowbridge [135] in their comparison between an acoustic Doppler velocimeter (ADV) and an LDV. Studying a very similar channel flow, they found that the ADV underestimated the horizontal velocity by 1% relative to the LDV measurements. This value is the same as what was observed here between PIV and LDV. Voulgaris and Trowbridge attributed their observed bias to sensor height mismatch. However, another possibility may explain the bias. McLaughlin and Tiederman [98] first recognized that the LDV individual realization sampling scheme is biased toward higher velocities because the volume of fluid, and therefore the number of seed particles, swept through the measurement volume is proportional to the fluid velocity itself. This bias has typically been a concern in very high velocity flows where the turbulent intensities are large. In the flow studied here (and in that of Voulgaris and Trowbridge) the bias would be expected to be very small. However, a precise comparison is desired in this case and such a bias may play a role.

Several post-experiment correction methods have been proposed to account for this velocity bias (*e.g.*, Herrin and Dutton [53]). A common approach is to correct mean flow field quantities using a weighting factor. To calculate an average quantity of an arbitrary velocity statistic x , the following equation is of use:

$$\bar{x} = \frac{\sum_i x_i w_i}{\sum_i w_i}, \quad (4.21)$$

where w_i is a weighting factor and the summations are taken over the entire record. The uncorrected, biased, individual realization case corresponds to $w_i \equiv 1$. Numerous forms of

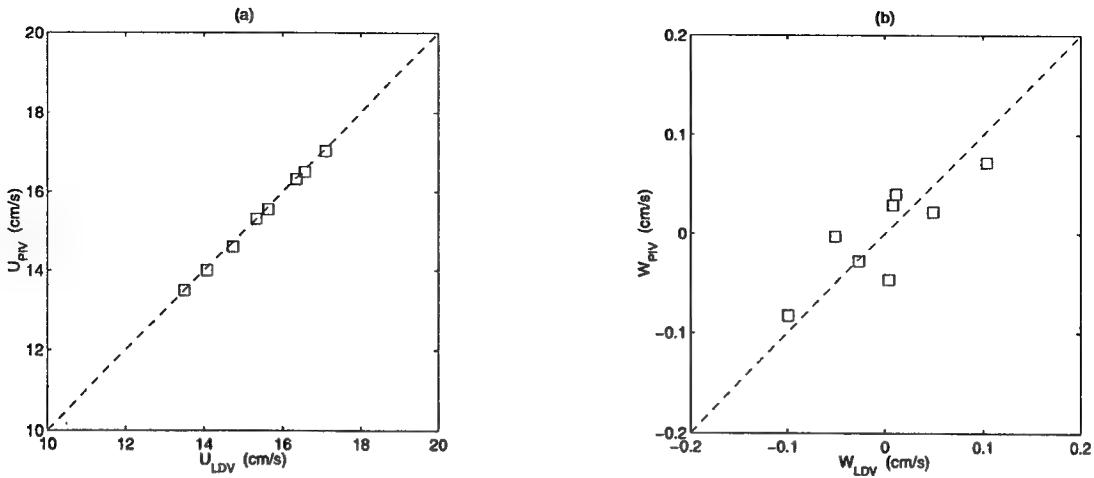


Figure 4-8: Comparison between mean velocity components measured by LDV (bias-corrected) and PIV. The dashed line represents the 1:1 relationship.

the weighting factor have been proposed: inverse velocity magnitude, particle interarrival time, and particle residence time are common. For the LDV data collected during this group of validation experiments, the particle interarrival time was selected as the weighting factor since in subsequent experiments the velocity magnitude was not completely known (only a single component was able to be measured with the LDV). The interarrival time weighting factor has been found to perform favorably in comparison analyses (Herrin and Dutton [53]).

Figure 4-8 shows the LDV bias-corrected comparison results for the mean horizontal and vertical velocities measured by the LDV and PIV systems. As stated earlier, the agreement is excellent. For the horizontal velocities, the mean error with the correction performed is only 0.06 cm/s. (The term “bias” is not used since some PIV measurements are higher than that of the LDV.) The RMS error was 0.07 cm/s. Incidentally, the mean error after using the inverse velocity magnitude weighting factor correction was 0.05 cm/s and the RMS error was 0.06 cm/s. These errors are only $\approx 0.4\%$ relative to the LDV measurement. This level of error could conceivably derive from the contributions from the various PIV sources discussed above. From this very encouraging agreement, it is concluded that the PIV system is capable of measuring mean flow velocities under favorable conditions with

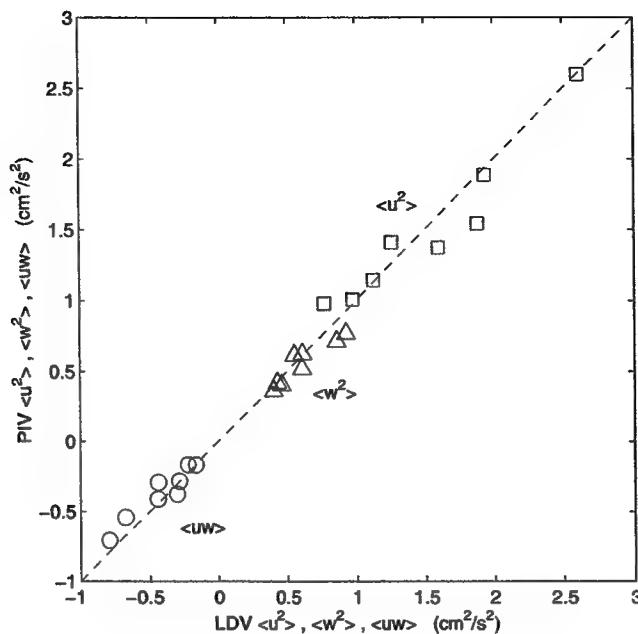


Figure 4-9: Comparison between turbulent flow statistics measured by LDV (bias-corrected) and PIV. Variance of streamwise velocity fluctuations (squares), variance of vertical velocity fluctuations (triangles), and velocity fluctuation covariance (circles) shown. The dashed line represents the 1:1 relationship.

errors on the order of 1%.

In addition to mean flow analysis, turbulent flow statistics were also evaluated for the two instruments. These results, which reflect the LDV bias correction, are shown in figure 4-9. Good agreement exists between the LDV and PIV measurements of the variances of the two velocity components as well as their covariance. Differences are most likely due to the fact that the two measurements were not coincident in space, and the flow sampled by the PIV system may not be able to be considered statistically equivalent to that sampled by the LDV over the relatively short measurement period.

Finally, figure 4-10 shows the result of the 10 minute times series experiment with the two systems. The height above bottom was 2 cm. The LDV data rate averaged 44.02 Hz for the length of the record. This data was resampled using linear interpolation to 44 Hz for processing. The PIV data was processed with 32×32 subimages (3 mm square). The agreement for both velocity spectra is excellent over the coincident frequency range. A noise

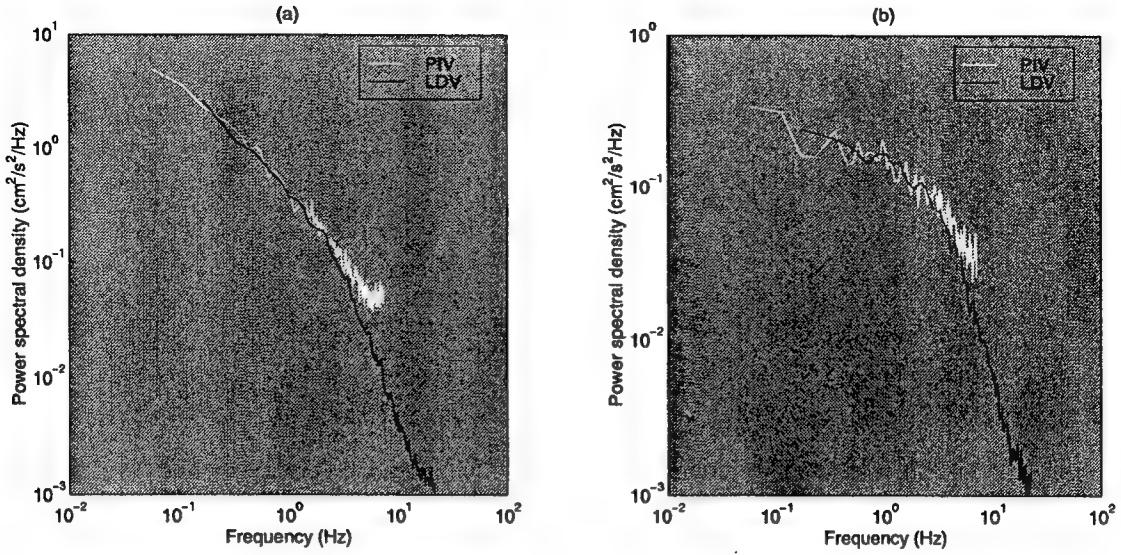


Figure 4-10: Velocity spectra as measured by LDV and PIV in the phase 1 channel flow, 2 cm above the bottom. (a) Streamwise velocity component, and (b) vertical velocity component.

floor for the PIV measurement seems to appear near 4–5 Hz. Slightly before this frequency, the PIV begins to show higher spectral energy than the LDV result, which may be an effect of measurement noise contaminating the spectra. Aside from that small discrepancy, the PIV measurement is able to provide good, low frequency spectral estimates.

4.5 Validation of the technique: Phase 2

In phase 1 of the validation, the core PIV system, which consisted of the hardware implementation and the velocity extraction algorithms, was shown to accurately measure bulk flow as compared to a laser Doppler velocimeter. In phase 2 of this validation, the PIV system was used in the proposed free-surface measurement configuration as described in section 4.2, and velocities at the surface were measured. Again, the LDV system was used as an independent measurement. The goal of this comparison was simply to determine if the new arrangement used in the surface imaging technique introduced any errors to the measurement. Before considering a sheared free-surface flow, a simple channel flow with a vertical streamwise velocity profile was examined. Since the flow velocity should be constant

in the near-surface region, errors incurred in the PIV measurement would be reflective of the proposed implementation, and not from any effects of imaging particles at depth.

4.5.1 Experimental arrangement and procedure

Experiments again were conducted in the paddle-driven racetrack flume. The water depth was 18 cm and the nominal flow speed was 4.5 cm/s. The surface flow was measured at a mid-channel location 6 m downstream of the start of the linear raceway section. The LDV was set up as before, but with the following modifications. In order for the LDV to make a free-surface measurement, the head unit was put at a rising angle of 1° relative to the plane of the interface. This angle has no effect on the horizontal measurement. Additionally, only the horizontal measurement axis could be used in this surface configuration since the upper beam necessary for the vertical measurement was directed through air. Vertical positioning of the measurement volume was done both visually and through interpretation of the LDV data rates. The former involved vertically stepping the LDV and visually determining when the point of internal reflection occurred at the intersection point of the two beams. In addition, monitoring the LDV data rate provided further evidence of a surface measurement. This was due to the fact that the surface contained a large number of floating (PIV) seed particles. Since the seeding density was greater at the surface than in the bulk, a sharp jump in the data rate was observed when the measurement volume was stepped vertically to the interface; above this point, the data rate fell off abruptly, indicating a measurement volume that was internally reflected. Using these two approaches, the LDV measurement volume was estimated to measure the flow within 200 μm of the free surface.

The PIV system was set up to illuminate and image a small patch of the free surface that included the location of the LDV measurement volume. The flow was seeded several meters upstream with the fluorescent particles described in section 4.2.2, and the camera was filtered using the 560 nm optical bandpass filter. The complete setup was like that depicted in figure 4-11. An image from the PIV camera is given in figure 4-12 that shows the camera field of view relative to the LDV beams and measurement volume. The precise *xy*-location of the LDV measurement volume was determined by analysis of the image in figure 4-12. A first-order least-squares curve-fit was applied to each beam image, and the intersection

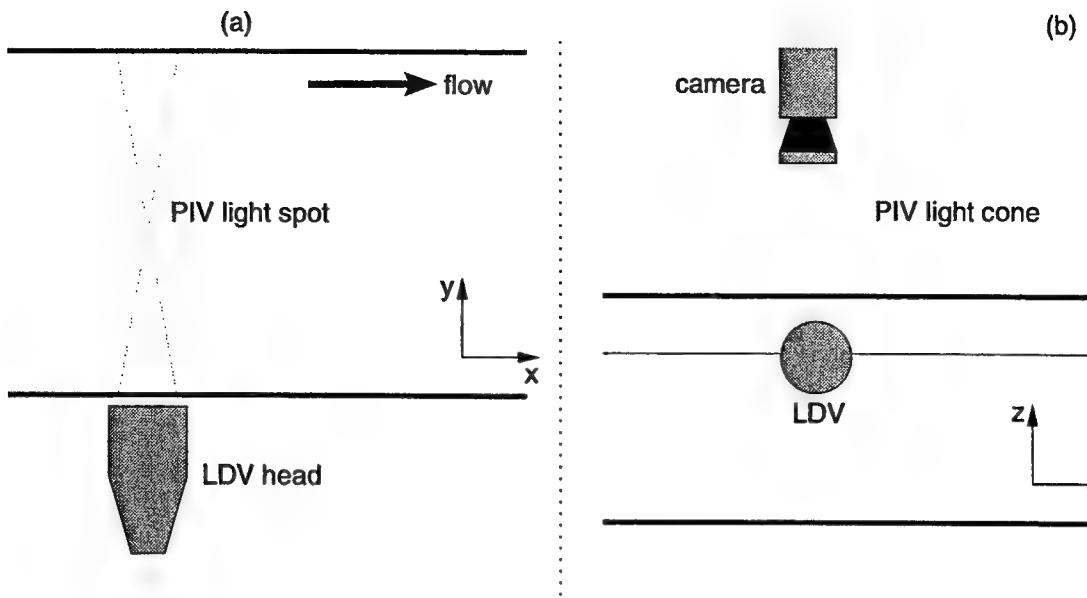


Figure 4-11: Experimental setup for phase 2 comparison between LDV and surface PIV systems studying a channel surface flow. (a) Plan view of channel. (b) Side view. Nd:YAG laser was situated above the flume channel. Not to scale.

of the two resulting lines was used as the location of the measurement volume. The PIV image acquisition was configured such that full resolution images of dimensions 1008×1018 were acquired in continuous mode. The corresponding physical domain was 11.1 cm in the streamwise direction by 11.0 cm in the cross-stream direction. The time separation between laser pulses was 12 ms; streamwise particle image displacements were on the order of 5 px. Because the full spatial resolution of the camera was utilized, only a short sequence of images could be collected by the image acquisition system. Using all available memory, 150 images were collected, yielding almost 5.2 seconds of PIV velocity data at 14.5 Hz. Ten seconds of LDV data was collected, and the two systems were synchronized to begin collecting data at the same moment. Timing synchronization errors were estimated to be less than 500 ms. Five experimental runs were conducted for the comparison.

4.5.2 Analysis and results

In order to ascertain the nature of the velocity profile near the free surface, LDV measurements of the velocity at $z = 0$ and $z = -4$ mm were made (the vertical origin is now the free

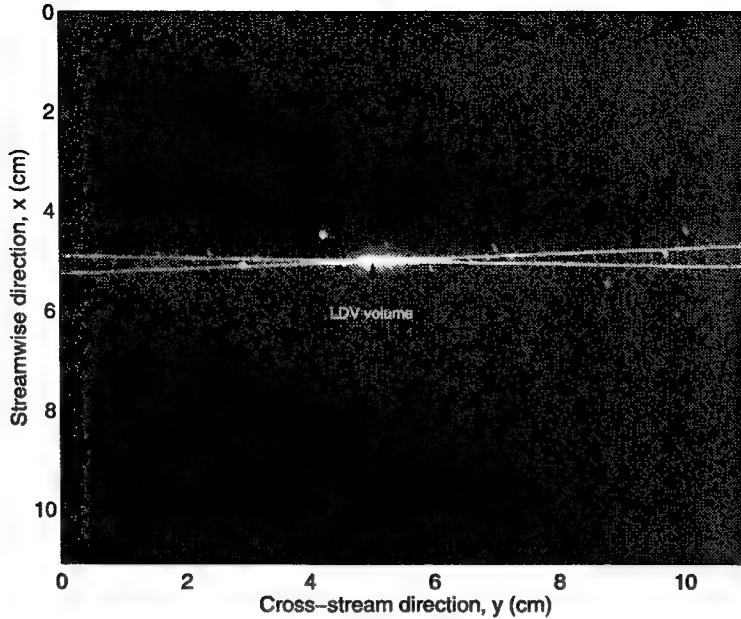


Figure 4-12: PIV camera image showing camera field of view with LDV beams and measurement volume. The view is from above, normal to the water surface. Note, the optical filter is not in use. The LDV beams are traveling right to left. The beams to the left of the measurement volume have been internally reflected by the air–water interface. Axis origins are arbitrary.

surface, with positive still upward). A 5 minute time series was taken at $z \approx 0$, followed by a 5 minute series at $z = -4$ mm. This pair of measurements was repeated a second time. The first run yielded a surface mean velocity of 4.741 ± 0.005 cm/s, and a mean velocity at 4 mm depth of 4.765 ± 0.009 cm/s (range indicates 95% confidence level). The second run yielded a surface mean velocity of 4.798 ± 0.008 cm/s, and a mean velocity at 4 mm depth of 4.777 ± 0.006 cm/s. These results confirmed a velocity profile at the free surface that was vertical. Thus, depth-averaging errors of the PIV technique should not come into play.

The LDV data rate averaged over the five runs was 59.0 Hz. Only the first 5.2 seconds of each LDV record was used in the comparisons since this was the segment that was coincident with the PIV measurement. The image data were processed for both DPIV and particle tracking velocimetry (PTV). For the DPIV processing, the hybrid correlation method was used with subimages of dimensions $32 \text{ px} \times 32 \text{ px}$. This yielded a measurement area that was $3.5 \text{ mm} \times 3.5 \text{ mm}$ in the xy -plane. Velocity vectors were computed every

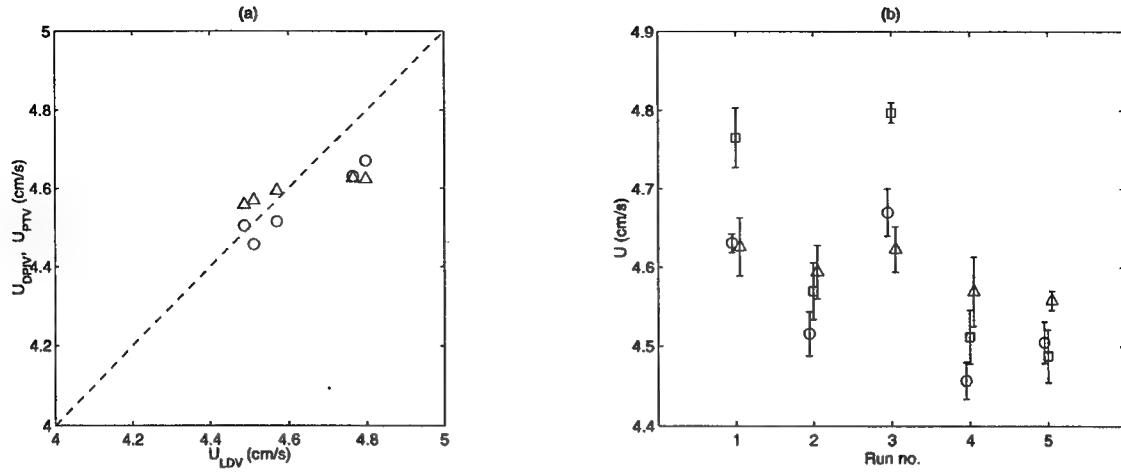


Figure 4-13: Mean surface velocities as measured by LDV (squares), DPIV (circles), and PTV (triangles) techniques for a basic channel flow. (a) Comparison between image velocimetry methods and LDV method. The dashed line represents the 1:1 relationship. (b) Mean velocities from all three techniques showing measurement variability. The vertical bars indicate 95% confidence intervals.

6 px (0.7 mm) in both dimensions. Again, outlier removal was enabled, but no further post-processing was performed. For each velocity field, the streamwise velocity at the location of the LDV measurement volume was determined by nearest neighbor interpolation. A 14.5 Hz time series of this measurement was used for comparison with the simultaneous LDV measurement. The PTV processing yielded irregularly gridded particle velocity measurements for each image pair. The vectors in each velocity field were used to interpolate a velocity at the LDV measurement volume location. This value formed the PTV time series, also at 14.5 Hz.

The results from these experiments are summarized in figure 4-13. The LDV results were bias-corrected using the particle interarrival time weighting method. Figure 4-13a shows how the results from the PIV and PTV methods compare to those of the LDV. The mean velocity error between the LDV and PIV was 0.07 cm/s; the RMS error was 0.09 cm/s. Between the LDV and PTV, the mean error was 0.03 cm/s and the RMS was 0.11 cm/s. The PIV results are similar to those found in the bulk flow experiment, and the PTV processing appears to yield comparable results. Relative to the LDV, the errors range from roughly 0.6–2.4%. The increase in relative error compared to the results from the bulk

flow measurements may be due to the more challenging measurement location at the free surface where aspects like particle seeding are less favorable. Shown in figure 4-13b are the mean velocities measured by the three techniques and the respective confidence intervals for each. The confidence results imply that, over all five runs, PIV yielded a slightly better estimate of the mean streamwise velocity; the LDV and PTV estimates of the mean were virtually identical in confidence. Based on the results from these experiments, it does not appear that the surface PIV configuration modifies the performance of the standard PIV system.

4.6 Validation of the technique: Phase 3

In the final phase of the validation, a flow with near-surface shear was studied. In such a case, the flow velocity varies with depth, thus providing the true test for the proposed surface measurement technique. Validation experiments were performed on two occasions. For the first experiment, a second PIV system was used to measure the near-surface flow independently. In the second experiment, the familiar LDV system was used to measure the surface velocity. The first experiment was performed very early on, and involved certain experimental complications, preventing an ideal comparison and prompting the sequence of LDV comparison experiments. Irrespective of the complications with the PIV-PIV experiment, its results strongly support the validity of the surface PIV measurement technique developed in this thesis.

4.6.1 PIV-PIV comparison

Both comparisons, that using the second PIV system and that using the LDV system (described in the next section), were performed in the racetrack flume. However, a paddle-driven channel flow was not used. Instead, a wind-driven flow was employed to create the desired free-surface shear layer. To accomplish this, a small fan was positioned over the water surface just upstream of the measurement region (see figure 4-14). To better direct the airflow, the fan was fitted with a plenum. Additionally, the fan was modified such that its speed was adjustable. The airflow rate was adjusted such that visible waves were not

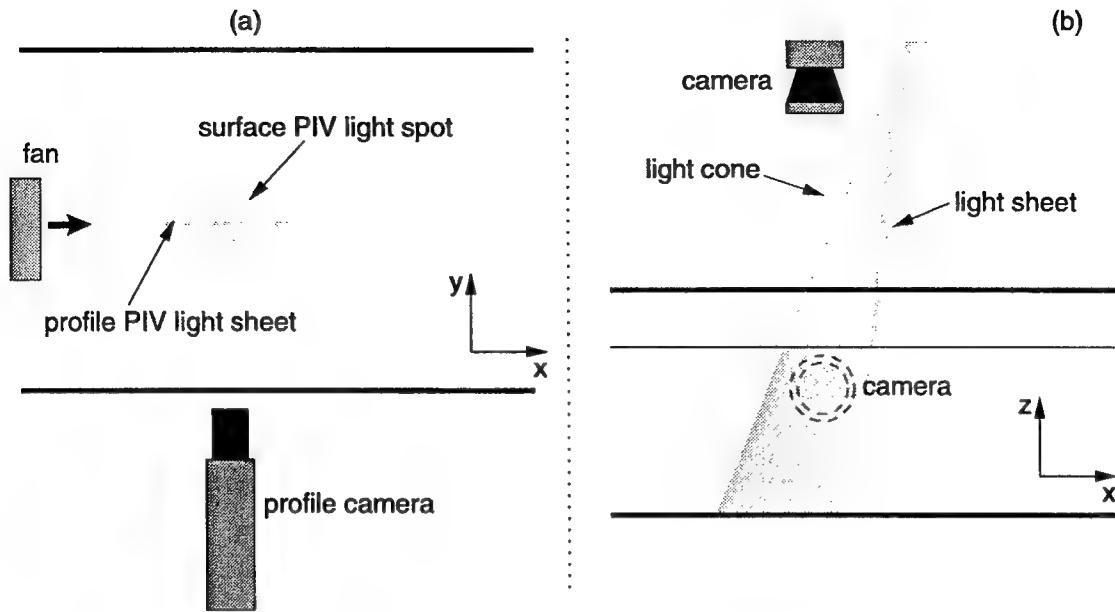


Figure 4-14: Experimental setup for phase 3 comparison. Shown is the arrangement for the PIV-PIV experiment. (a) Plan view. (b) Side view. Not to scale.

produced at the surface. For both experiments, the approximate distance between the fan and the measurement location was 50 cm.

The surface PIV measurement system was arranged as described earlier in section 4.5.1. Again, the flow was seeded with fluorescent particles upstream and the camera optically filtered. To acquire an independent measurement of the surface velocity, a second PIV system was deployed with the intent of measuring the velocity profile right up to the air-water interface. This system was composed of a frequency-doubled Nd:YLF pulsed laser, a CCD camera, and supporting video acquisition hardware and software. The laser was run at 20 kHz, thus appearing continuous for PIV purposes, and emitted green light at 523.5 nm in the 3 W range. Optics were arranged that directed the laser beam from the head unit (located above the channel) through a Powell lens, which formed the beam into a 1–2 mm thick light sheet. The light sheet was directed downward, through the air–water interface, illuminating a vertical xz -plane of the flow. This plane was located coincidently within the xy -plane imaged by the surface PIV camera. Ambient particles in the water were used as flow tracers for the profiling measurement. The flow was slow enough so as not to require

shuttering of the Nd:YLF laser. Therefore, each video image was exposed for a full 33 ms.

The profile PIV camera was an analog CCD camera with resolution 1134 (C) \times 486 (R). The video capture board available for use was unable to retain the spatial resolution of the camera and digitized the images to a reduced resolution of 640×480 . Details of the complete system are given in McKenna [96]. The camera was set up alongside the channel test section and was oriented with a slight rising angle relative to the free-surface plane. This enabled the internal reflection of the subsurface to be imaged and provided a means to locate the position of the free surface. This angle impacted the scaling of the vertical velocity, which was not analyzed, and had no effect on the streamwise velocity. The camera field of view corresponded to $6.9 \text{ cm} \times 5.2 \text{ cm}$.

It originally was hoped that temporally coincident measurements from both systems could be made. However, it became apparent that the light cone from the surface PIV system produced an intense free-surface illumination that was imaged by the profiling PIV camera. Because of this strong brightness, the profiling camera was unable to visualize the flow very near the interface. Optically filtering this camera proved unsuccessful. As a result, the measurement strategy was altered. In lieu of simultaneous measurements, approximately 2.5 seconds of profile PIV data were collected, immediately followed by approximately 2.5 seconds of surface PIV data. Longer time records were not taken since the wind-driven flow was weakly unsteady.

Five runs were performed in all, but only the results from a run that showed the best image data quality will be presented. The profile measurements will be considered first. Image pairs were processed with direct spatial correlation using 8×8 subimages ($0.9 \text{ mm} \times 0.9 \text{ mm}$). The complete images were processed. This meant that the true flow field and the internally reflected flow field were considered as a single flow. Velocities were computed every 0.2 mm in the vertical and every 1.3 mm in the horizontal directions. For each velocity field, a mean profile was calculated by averaging the velocities in the streamwise direction. These mean profiles were averaged in time to yield a final space-time mean velocity profile, denoted $\langle U \rangle$.

Figure 4-15 shows an image taken from the sequence and a vector plot of the mean, temporally averaged, spatial velocity field for this run. The location of the free surface

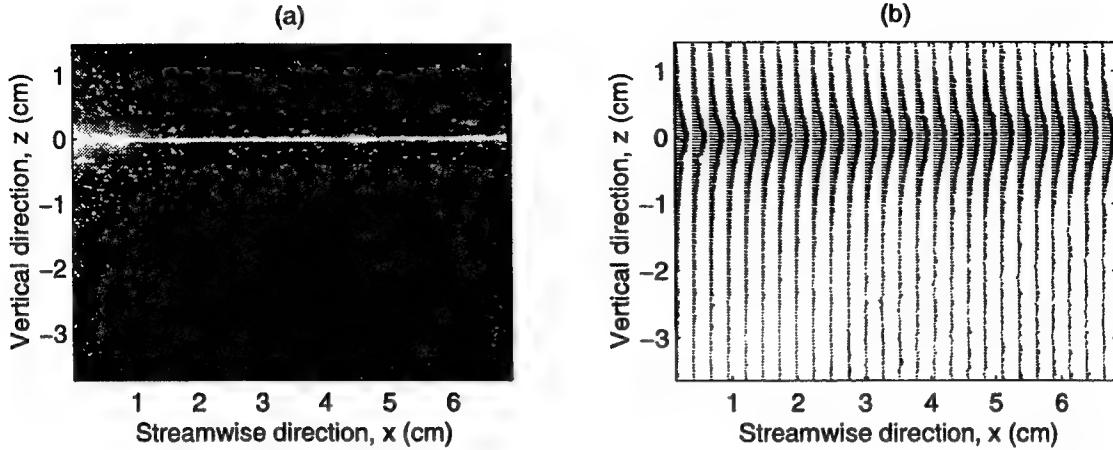


Figure 4-15: (a) Sample profiling PIV image. The free surface is seen as a thin bright band. The subsurface reflection is seen for $z > 0$. (b) Temporally averaged, mean velocity field result. Every other profile is shown for clarity. The horizontal axis origins are arbitrary.

was estimated by computing the center of the thin horizontal light strip representing the interface. Over the sequence, this strip measured 0.6–1.0 mm in thickness. If the actual interfacial boundary lied in the middle of this strip, measurements of the velocity profile would only be possible up to within 300 μm of the surface at best. Incorporating the finite size of the correlation subimage, this limit is more like 750 μm . This illustrates some of the difficulties eluded to in section 4.1 regarding accurate and precise near-surface profiling measurements using imaging velocimetry. In addition to obscuring the exact interfacial location, the surface light strip also impacted the DPIV correlation processing. Because the light strip remained a permanent feature in the images, it acted as a rigid band that did not move with the free-surface velocity. This resulted in a velocity bias toward zero for the correlations performed with subimages that included this strip.

Figure 4-16 contains the mean profile data result. The bias toward zero in the velocity as the measurement gets very near the interface is evident; the shape of the profile inflects back toward surface. Because these near-surface measurements were considered biased, a least-squares curve-fitting procedure was used to extrapolate the mean velocity profile $\langle U \rangle$ to the interface. Assuming a model of one-dimensional momentum diffusion into a

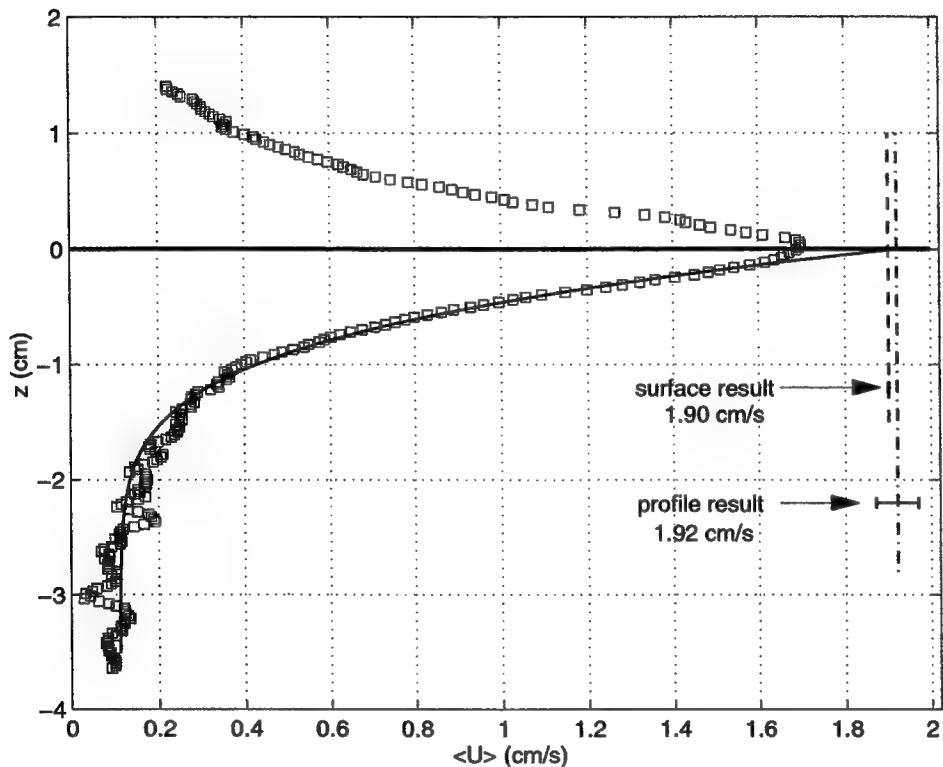


Figure 4-16: Result from PIV-PIV surface velocity comparison experiment. Shown is the mean measured velocity profile as actual data (squares) and fitted curve from equation 4.22 (solid line). The free surface is at $z = 0$. The data for $z > 0$ were obtained by processing the internally reflected image of the subsurface. The result from the profiling PIV measurement was an extrapolated surface velocity of 1.92 ± 0.05 cm/s, given by the dash-dot line. The surface PIV system yielded a final result of 1.90 ± 0.002 cm/s, given by the dashed line. The horizontal bars indicate the measurement 95% confidence intervals as discussed in the text.

semi-infinite domain, the expression determining the unsteady velocity profile is

$$\frac{U(z, t)}{K} = z \left[1 + \operatorname{erf} \left(\frac{z}{2\sqrt{\nu t}} \right) \right] + 2\sqrt{\frac{\nu t}{\pi}} e^{-z^2/4\nu t}, \quad (4.22)$$

where $K = \tau_0/\mu_{\text{water}}$, with τ_0 representing the surface stress due to the wind (White [142]). Ultimately, the profile assumes a steady shape due to the balance between the energy input to the system from the wind and the energy losses due to the fluid viscosity and stresses at the boundaries. A least-squares curve-fit of the form given by equation 4.22 was used to fit the velocity data for $z < -0.18$ cm. From this curve-fit, the extrapolated surface velocity was found to be 1.92 cm/s. To estimate the confidence in this estimate, the following procedure was used. When reducing each spatial velocity field to a single mean profile, the 95% confidence interval at each vertical position was computed. This vertical distribution of confidence intervals was averaged over all velocity field instances to yield a final vertical distribution, denoted $C_U(z)$. Two more curve-fits were produced from

$$\langle U \rangle_- = \langle U \rangle - C_U(z) \quad (4.23)$$

$$\langle U \rangle_+ = \langle U \rangle + C_U(z). \quad (4.24)$$

The extrapolated surface values resulting from these two fits were 1.87 cm/s and 1.97 cm/s, respectively. Therefore the final estimate was considered to be 1.92 ± 0.05 cm/s. This is shown in figure 4-16. Incidentally, an estimate of the surface stress can be found using the fit parameters. Based on the fit, the mean stress at the interface was approximately 2.5 mPa ($u_* = 0.16$ cm/s).

In the surface PIV analysis, images were processed with direct spatial correlation using 32×32 subimages (2.5 mm \times 2.5 mm). Velocities were computed every 0.6 mm in both directions. For each velocity field, the flow was considered homogeneous along the streamwise direction. The variability along this direction was estimated as being on the order of 0.5 mm/s over the measurement region. The velocities were collapsed in that dimension, and a linear least-squares fit of $u = u(y)$ was calculated. The standard deviation of the velocities that were used to construct each fit was approximately 0.03 cm/s, on average. Using the result of the fit, the velocity was evaluated at the exact y position of the profiling

PIV light sheet. The final surface velocity estimate was the time-average of this value: 1.90 ± 0.002 cm/s. The range indicates the 95% confidence level for the time-averaged estimate of the mean. Figure 4-16 shows how this result compares with that of the profile PIV result.

Based on these results, it is concluded that the surface PIV technique is accurately measuring the velocity field at the interface, and effects of depth-integration are insignificant and can be disregarded. A worst case estimate can be considered by assuming the surface flow velocity was $1.92+0.05 = 1.97$ cm/s. If this were the case, the surface PIV measurement would have underestimated the true value by 0.07 cm/s. Using the fitted mean profile, we can determine at what depth such a velocity reduction would be observed. Taking into account a depth average, this would imply a measurement over a depth of approximately 550 μ m. A best case estimate would imply a measurement *at* the surface. While the result from this comparison was extremely favorable, certain aspects of the experiment were imperfect. First, the measurements were not precisely simultaneous. Second, the PIV profile measurement was unable to resolve the surface velocity directly and required a theoretical extrapolation to produce an estimate at the interface. With these caveats in mind, a second experiment was conducted using the LDV system employed in the previous validation experiments.

4.6.2 LDV-PIV comparison

The set up for this experiment was a combination of the LDV arrangement already described in section 4.5.1 and the wind-generated flow discussed in section 4.6.1. The surface PIV system imaged and measured the wind-driven flow in a small surface patch, and the LDV measured the surface flow in a measurement volume contained within the PIV patch. The water depth was 18 cm. The PIV processing parameters were: hybrid correlation, field of view = 11 cm square, 32×32 subimage size (3.5 mm \times 3.5 mm), measurement grid spacing = 0.7 mm in both directions, $\Delta t = 16$ ms, and nominal particle image displacement = 3–4 px. Particle tracking also was performed. Ten runs were conducted. One of the runs suffered from poor PIV seeding and therefore was not analyzed. The sequence length was limited to 150 images, or approximately 5 seconds of velocity data at 14.5 Hz. The

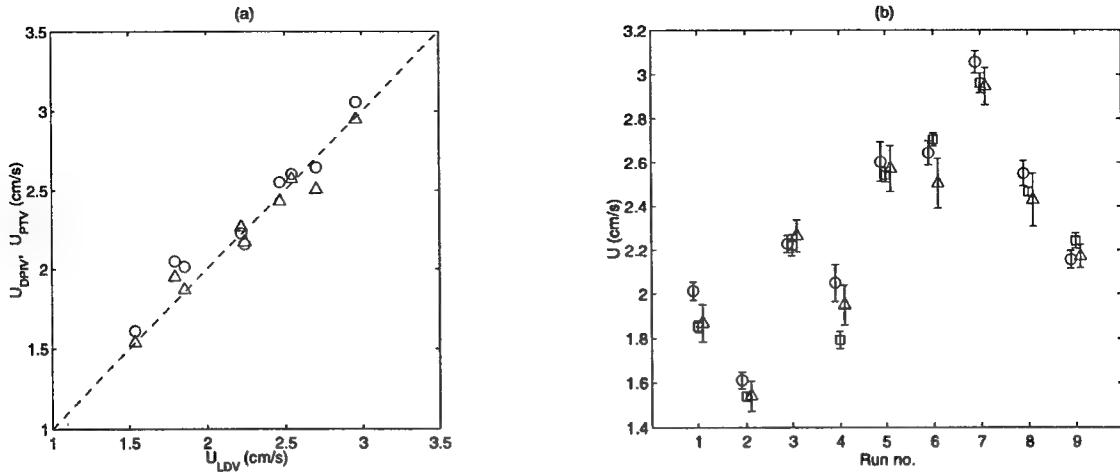


Figure 4-17: Mean surface velocities as measured by LDV (squares), DPIV (circles), and PTV (triangles) techniques for a wind-driven flow. (a) Comparison between image velocimetry methods and LDV method. The dashed line represents the 1:1 relationship. (b) Mean velocities from all three techniques showing measurement variability. The vertical bars indicate 95% confidence intervals.

LDV data rate averaged 253 Hz over the nine studied runs. Time series from both image velocimetry methods were formed in the same manner as described in section 4.5.2

Figure 4-17 presents the results of the comparison in terms of mean surface velocities. Figure 4-17a shows the relationship between all three measurement methods, and figure 4-17b indicates the respective measurement confidences. The mean velocity error between the LDV and PIV was -0.07 cm/s (PIV higher than LDV); the RMS error was 0.12 cm/s. Between the LDV and PTV, the mean error was 0.01 cm/s and the RMS was 0.09 cm/s. While the absolute errors encountered here are very similar to those found in the previous two LDV-PIV experiments, the relative errors are increased—roughly 0.4–5.3%. The fact that in this experiment, the PIV results were, in general, higher than those of the LDV is encouraging. This might indicate that PIV systematic errors do not necessarily imply underestimation. Therefore, based on these results we may conclude that the surface PIV measurement can accurately quantify the flow field at a flat air–water interface.

In all the experiments presented in this chapter, the PIV results have been quite remarkable. The agreement with the LDV system and the surface PIV system can be considered excellent. To provide a sense of this, figure 4-18 is included, which shows a snapshot of two

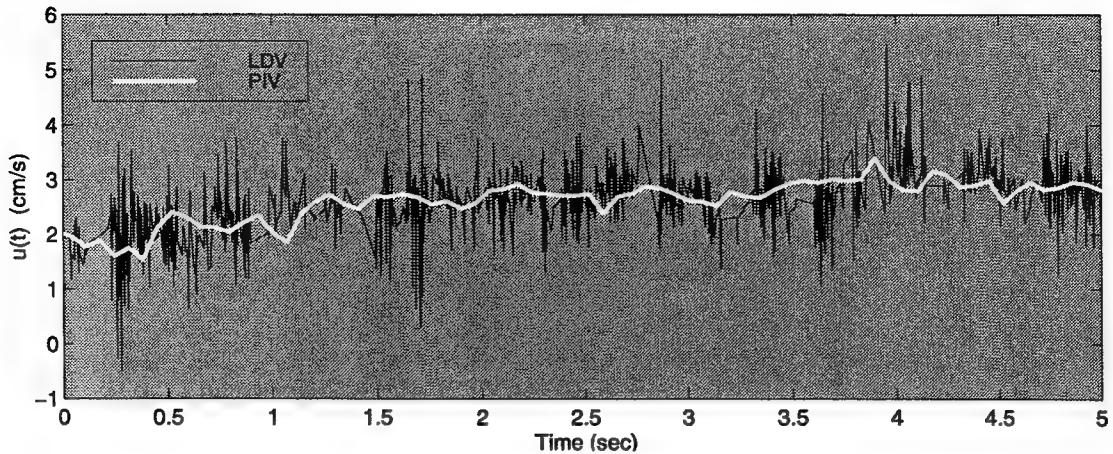


Figure 4-18: LDV and PIV time series $\tilde{u}(t)$ of wind-driven surface velocity for run no. 4. The average LDV data rate was 292 Hz. The PIV data rate was fixed at 14.5 Hz.

time series from the wind-driven LDV-PIV experiment. The LDV data rate was approximately 20 times that of the PIV rate for this particular run (no. 4). The LDV appears somewhat noisy, which also was revealed in the measurement standard deviation. This variability is likely due to the fact that the measurement volume was right at the free surface, and intermittently may have been internally reflected due to fine free-surface deformations.

4.7 Summary

In this chapter, the surface PIV measurement technique developed as part of this thesis work has been described and subjected to a series of strict validation experiments. While the technique itself is relatively straightforward, explaining its success is more involved. Theoretical considerations indicated that the presumption that seed particles at depth do not contribute to the measurement may be explained by small particle scattering theory. It was found that the intensity of scattered light from a particle at the free surface is many times greater than a particle in the bulk fluid. Visual observations and a simple experiment have contributed to verifying this behavior.

The experimental validation of this surface measurement technique was non-trivial. The challenge involved stems from the fact that extremely small discrepancies in velocity are

crucial to assessing the success of the method. In addition, obtaining a *definitive* measurement to use for comparison is always open to question, as a perfect measurement does not presently exist. This series of validation experiments, considered collectively, are believed to show that the proposed surface PIV measurement technique can successfully measure the flow field *at* a free surface, with possible depth effects no greater than a couple hundred microns. Processing the image data using PIV, and when possible PTV, revealed small errors on the order of 1-3%. These errors are within the probable error limits of the PIV technique in general, and could, of course, be partially attributable to the LDV as well. An advantage of the specific implementation of this surface measurement technique is its potential extension to deforming free-surface flows. The light cone illumination strategy, coupled with the optical filtering of the camera, allows this technique to image water surfaces exhibiting small slopes and small displacements from equilibrium. This is the subject of the next chapter.

Chapter 5

Free-surface PIV: application to plane progressive water waves

5.1 Introduction

This chapter is an extension of chapter 4: the capabilities of the surface PIV technique were further tested by applying this flow measurement method to the problem of surface water waves. Mechanically generated, plane, progressive, low amplitude waves were studied in a short linear wave tank. Independent and simultaneous measurements of one-dimensional wave slope at a single location in the tank were used to provide comparison results for the PIV.

5.2 Experimental arrangement and procedure

5.2.1 Wave tank and wavemaker

The progressive wave experiments were conducted in a 240 cm long linear wave tank. The wave channel was 40 cm across and 15 cm deep. The tank is constructed entirely from 1.27 cm thick clear polycarbonate sheet. Details of the tank can be found in Mass [91]. In all experiments, the water depth, h , was 10 cm. A paddle-type wavemaker was integrated into the tank (see figure 5-1). The paddle was a 6.4 mm thick polycarbonate plate that spanned the tank width and was positioned roughly 30 cm from one end of the tank. The

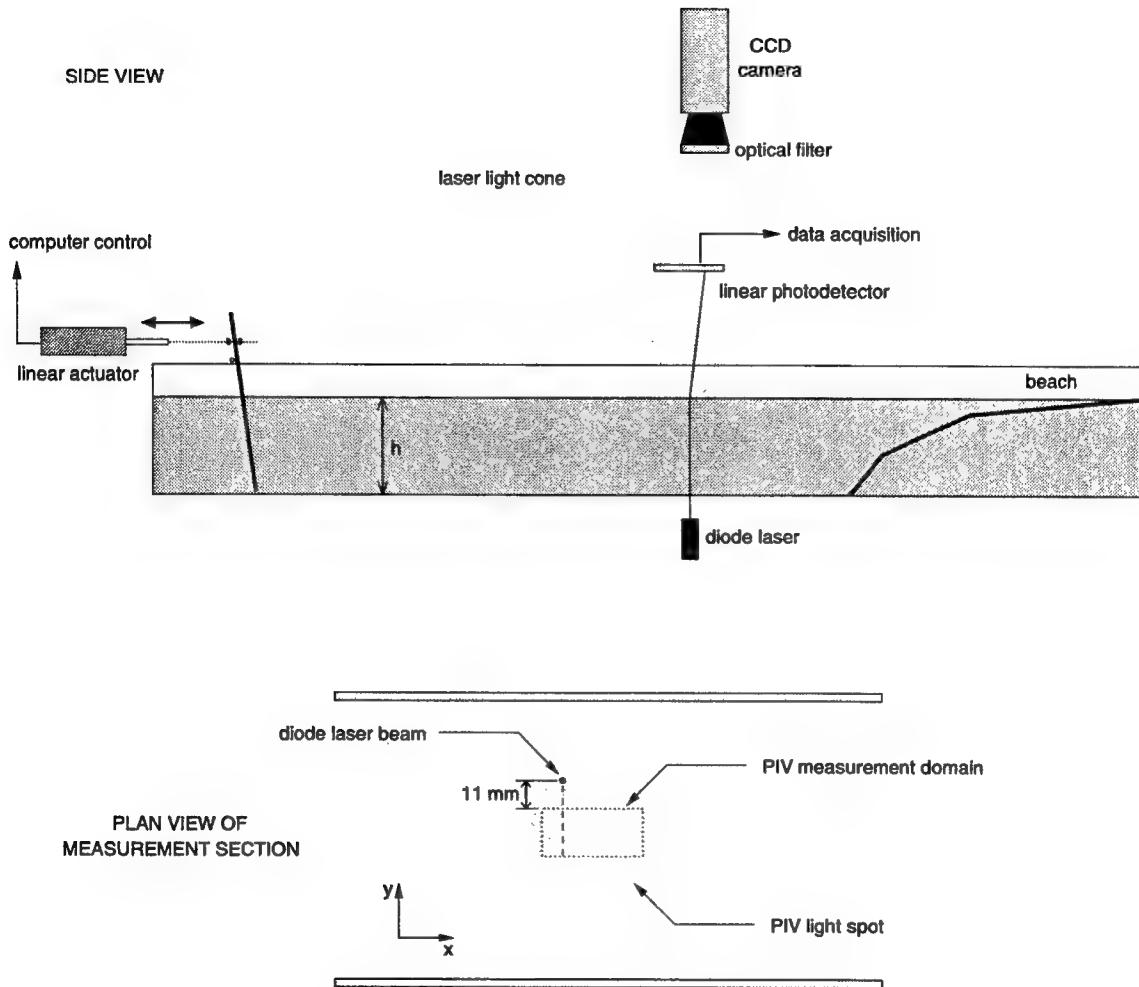


Figure 5-1: Diagram of linear wave tank setup. In the upper sketch, a side view of the wave tank is shown. The camera and photodetector are actually in slightly different planes so as not to interfere with each other. While not apparent from this sketch, the laser light cone was directed in a way that avoided the photodetector. The unrealistic refraction of the diode laser beam is shown for illustration only. In the lower sketch, the plan view of the measurement region is shown. The light spot was nominally 10 cm in diameter, and the measurement domain was roughly $8\text{ cm} \times 4\text{ cm}$. The wave slope measurement made with the diode laser was located about 1 cm away from the edge of the PIV measurement region. The short dashed line indicates the slice of the PIV velocity measurement region that was used in the comparison analysis. Neither sketch is to scale.

paddle was made to pivot about two points on top of the tank side walls. A threaded rod connected the paddle to a computer-controlled linear actuator. Sinusoidal paddle motions with frequencies in the range of 1–7 Hz were explored. The paddle excursion was adjusted in order to manipulate the wave amplitude. Wave amplitudes were kept small (on the order of millimeters) in order to allow for linear wave theory analysis. Both the linear theory assumption and the impact of finite tank depth are addressed in section 5.3.

Experiments were performed using spring water. However, cleanliness was not a requirement. In fact, some adventitious surfactant material was actually deemed desirable because it damped unwanted high-frequency waves that arose from slight vibration in the wave generation system. A single measurement of surface tension was made either before or after a group of data runs.

5.2.2 Surface PIV setup

The surface PIV system was setup as described in chapter 4. The laser light cone was directed from above at an unobtrusive angle toward the free surface, and the CCD camera was centered above the tank channel, focused on the light spot (figure 5-1). The light spot was nominally 10 cm in diameter. As usual, the water was seeded with the acrylic fluorescent particles. The image acquisition system was configured to sample a 768 px \times 384 px image and was run in continuous mode, yielding image pairs at 14.5 Hz. The resulting measurement domain was roughly 8 cm \times 4 cm. For a single run, a sequence of 672 images was collected. This provided 336 velocity fields over a span of about 23 seconds. The processing of the PIV data is discussed in section 5.3. The location of the measurement region was approximately 100 cm downstream of the wavemaker.

5.2.3 Wave slope measurement

In order to have an independent measure of the wave properties for comparison to the surface PIV velocity results, a wave slope point-measurement was implemented. Ideally, a wave height sensor would have been used since the wave orbital velocity field at the surface, $u(x, y, z = 0, t)$, is directly proportional to the wave amplitude ($u \sim A\omega$), and the radial wave frequency ω was easily measurable. However, a sensor capable of measuring

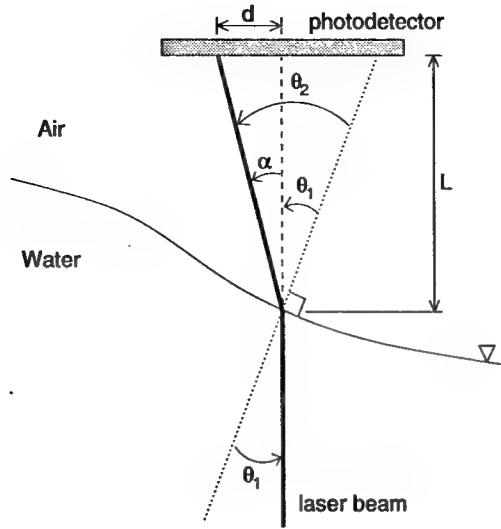


Figure 5-2: Diagram of wave slope measurement optics. The offset d , due to a finite wave slope, is related to the voltage output of the photodetector. The wave slope and d are related through equation 5.1. Not to scale.

the small wave heights encountered in these experiments with the necessary accuracy and precision was not readily available. Additionally, it is an invasive measurement. Wave slope is proportionally related to the wave amplitude by the wavenumber, $\kappa = 2\pi/\lambda$, as κA , with λ as the wavelength. Therefore, with knowledge of wave slope and the wavenumber (through a dispersion relation, if the wave frequency is known), the wave amplitude can be computed and related to the wave velocity field as above.

A non-invasive, optical method was used to measure one-dimensional wave slope at a single location in the wave field. The method is extremely straightforward and relies only on a low-power laser source and a linear photodetector. Here, a 3 mW diode laser (670 nm) was used. The laser was located below the wave tank, near center-channel, such that the beam was normal to the free surface. Above the tank, a commercially available, linear (one-dimensional) photodetector measured the position (in the x -direction) of the refracted diode laser beam as the wave field propagated by (see figure 5-2). The output of the photodetector was a voltage signal related to the position of the centroid of a small light spot traveling over its 2.5 mm \times 34 mm photosensitive area. This relationship was determined through a separate calibration experiment. Therefore, if the distance between the free surface and

the detector is known, the deviation of the beam position from its static location can be used to infer the slope of the surface. If the offset is denoted d and the distance between the free surface and the photodetector is L , the slope is related to the photodetector signal through Snell's Law (equation 4.3) as

$$\tan \theta_1 = \frac{n_2 \sin \alpha}{n_1 - n_2 \cos \alpha}, \quad (5.1)$$

where $n_1 = 1.331$ and $n_2 = 1.000$ are the indices of refraction for water and air, respectively, at the diode laser wavelength. The angle $\alpha = \text{atan}(d/L)$, and the slope of the interface is equivalent to $\tan \theta_1$. For these measurements, the signal from the detector was passed through an amplification circuit and then sampled by a digitizing oscilloscope at 500 Hz. A 40 second record was collected. The synchronization of the start times for the PIV image acquisition and the wave slope sampling was estimated to be within 500 ms.

5.3 Results and analysis

The data presented here are from two groups of wave experiments, each conducted on a separate occasion. Two runs from the first group (labeled 1-1 and 1-2) and four from the second (labeled 2-1, 2-2, 2-3, and 2-4) have been included. These runs were selected for analysis after qualitatively examining the characteristics of all collected wave slope records. Runs where the waves deviated noticeably from a pure sinusoid or were contaminated with an additional wave component were discarded. Therefore, the six data sets analyzed were deemed most amenable to standard linear wave theory analysis.

Three aspects of the wave flow were investigated: wave frequency, wavelength, and wave surface velocity. The analysis for each is discussed in the following three sections. Universal to all three was the processing of the PIV data. In all six cases, the hybrid correlation algorithm was employed. Additionally, vectors were computed every 16 px (1.7 mm) in both directions for all data sets. The only processing difference was that the group 1 runs were processed with 32 px \times 32 px subimages (3.4 mm square) and the group 2 runs were processed with 64 px \times 64 px subimages (6.7 mm square). The reason for the difference was due to increased particle displacements for the group 2 image pairs. Outlier removal

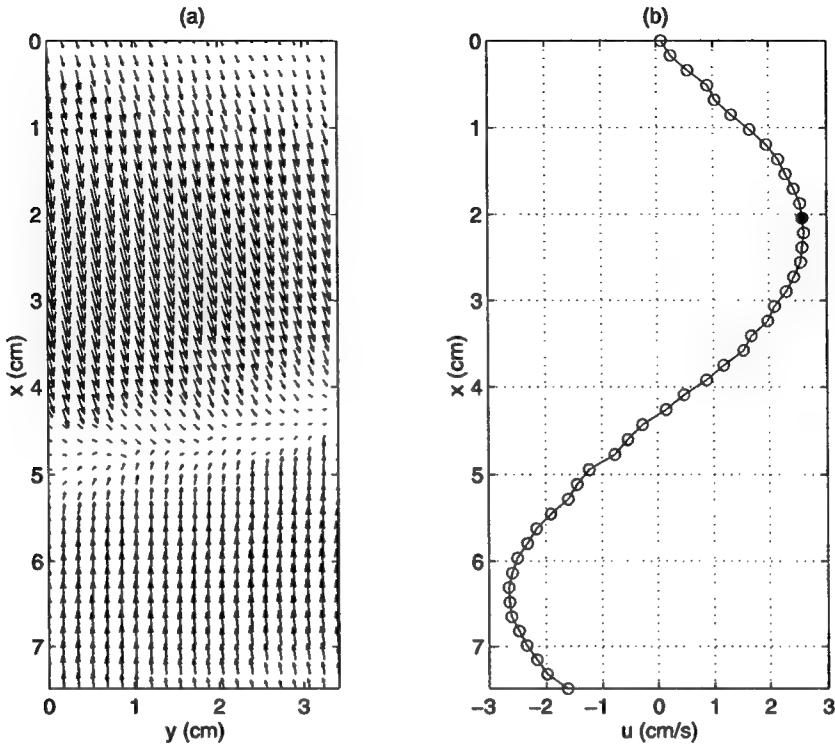


Figure 5-3: Sample wave surface velocity field. Result is from run 1-1. (a) Velocity vector field, and (b) cross-wave mean of the wave surface velocity, $u(x, z = 0)$. The solid marker indicates the sampled location used to form the single-point time series $u(t)$. The wave propagation direction is $+x$ (top to bottom). The x and y axes are arbitrary.

was enabled, but no further post-processing was performed.

To facilitate comparison with the single-point wave slope measurement, the PIV velocity field time record was reduced to a single-point time series. This was accomplished by first averaging each velocity field in the cross-wave direction yielding $u = u(x, z = 0)$ at each realization in time (see figure 5-3). Next, at each realization, the velocity at the equivalent wave slope measurement x -position in the tank (see figure 5-1) was collected to form $u = u(t)$ at the surface. This 14.5 Hz time series of velocity and the coincident portion of the wave slope times series were used to evaluate wave frequency and wave surface velocity. Wavelength was estimated by PIV using the 336 realizations of $u(x, z = 0)$, and by wave slope using a dispersion relation and the measured frequency. In this analysis, it will be posited that the results derived from the wave slope measurements are accurate,

Run	f_s (Hz)	f_v (Hz)	Error (%)
1-1	4.492 ± 0.003	4.491 ± 0.014	0.02
1-2	2.827 ± 0.001	2.826 ± 0.002	0.04
2-1	3.452 ± 0.005	3.453 ± 0.006	0.03
2-2	3.451 ± 0.003	3.454 ± 0.008	0.09
2-3	3.452 ± 0.002	3.451 ± 0.002	0.03
2-4	3.449 ± 0.003	3.457 ± 0.004	0.23

Table 5.1: Wave frequency results: f_s was determined from the wave slope record, and f_v was determined from the surface PIV velocity record. The error is defined as $|f_s - f_v|/f_s$.

and any discrepancies will be regarded as “error” attributed to the PIV measurement. This aspect is discussed in section 5.3.4.

5.3.1 Wave frequency

Both the PIV velocity time series and the wave slope time series were analyzed identically to extract wave frequency. Each record was binned into 5 second segments. A Fourier transform of each segment provided the frequency content of the wave signal within each 5 second bin. Around the peak (wave) frequency, a Gaussian curve-fit was used to improve the precision of the wave frequency estimate. The final wave frequency results are presented as the average of the frequencies obtained for each bin with the 95% confidence intervals indicated by plus/minus notation (table 5.1). The results computed from the wave slope data are denoted f_s and the results computed from the PIV data are denoted f_v . The excellent agreement between the two estimates of the wave frequency is evident. Discrepancies between the two are less than 0.1% in all cases but one. The surface PIV system has no difficulty is accurately measuring wave frequency. These results also indicate that the wave frequency (and the wavelength) can be considered constant over the sampling period.

5.3.2 Wavelength

To obtain a measurement of the wavelength using the wave slope data, it was necessary to invoke a dispersion relation that relates the wavelength, or wavenumber, to the wave frequency. For one-dimensional, plane progressive waves in deep water, linear inviscid theory

specifies

$$\omega^2 = g\kappa, \quad (5.2)$$

where $\omega = 2\pi f$. The validity of the assumption of linear waves is treated in section 5.3.3 and the assumption of deep water is investigated at the conclusion of this section. If surface tension is included in the free-surface boundary condition, the dispersion relation becomes that due to Kelvin [131]:

$$\omega^2 = \frac{\sigma\kappa^3}{\rho} + g\kappa. \quad (5.3)$$

Surface tension can be safely ignored in the dispersion relation for small wavenumbers, but does enter into the relation as the wavenumber increases. For the waves analyzed, the effect of surface tension on the wavenumber for a wave frequency of 2.8 Hz (lower frequency bound) is only 0.7%. For a frequency of 4.5 Hz (upper frequency bound), the effect on wavenumber is 4.3%. Therefore, throughout the analysis in this chapter, equation 5.3 is used to infer the wavenumber from the measured wave frequency obtained from the wave slope record.

A more direct measurement of the wavelength can be obtained with the PIV velocity field data. To accomplish this, each realization of $u(x, z = 0)$ was curve-fitted with a sinusoid. From each fit, an estimate of the wavelength was obtained. These estimated wavelengths were averaged and 95% confidence intervals computed. Due to the actual wavelengths studied and the size of the PIV measurement region, the comparison between estimates of wavelength from the two measurement systems was limited to a single data set: run 1-1. This run was at a high enough frequency such that roughly an entire wavelength was captured in the PIV measurement region. The amounts of wave phase contained in the measurement region for the lower frequency runs prohibited accurate curve-fits to the data. Using the two approaches just outlined, the estimated wavelength using the wave slope data was $\lambda_s = 8.05 \pm 0.01$ cm. The result using the PIV data was $\lambda_v = 8.32 \pm 0.04$ cm. The error, defined as $|\lambda_s - \lambda_v|/\lambda_s$, is 3.35%. Since the agreement between the frequency measurements was nearly exact, it appears that the PIV result may slightly overestimate the wavelength.

This discussion of wavelength concludes with a look at the effect of the tank finite depth.

Ignoring surface tension, for a fluid domain with depth h , the deep-water dispersion relation (equation 5.2) becomes

$$\omega^2 = g\kappa \tanh kh. \quad (5.4)$$

Thus, it is the product kh that determines the influence of finite depth on the relationship between frequency and wavelength. The above equation tends to the deep-water limit (equation 5.2) for $kh \gg 1$, although a condition of $kh > \pi/2$ is sufficient; the shallow-water limit obtains for $kh \ll 1$. Of the waves analyzed, the case most influenced by finite depth is run 1-2, where the measured frequency was approximately 2.8 Hz. For this wave, the wavenumber as computed using the deep-water expression is 0.321 cm^{-1} ; the result using the shallow-water expression is 0.322 cm^{-1} —an insignificant difference for this analysis.

The surface velocity field, examined in the next section, can also be influenced by finite depth through the modification of the wavenumber. For deep water, $u \sim A\omega$. In shallow water, the dependence is $u \sim gA\kappa/\omega$, where κ is determined by equation 5.4. However, for the waves studied, the impact of the finite depth on the wavenumbers is too small to have any appreciable effect on the surface velocity. Therefore, it can be concluded that finite depth effects are unimportant and that the use of the deep-water Kelvin dispersion relation is justified. Inclusion of surface tension will not affect this conclusion.

5.3.3 Wave surface velocity field

The final wave characteristic examined was the wave surface velocity, which is closely related to the wave amplitude, $u \sim A\omega$. As already mentioned, the wave frequency can be treated as constant, and therefore, a measurement of the wave amplitude, coupled with the measured wave frequency, can be used to calculate the wave orbital velocity at the free surface. This methodology was necessary in order to infer the wave surface velocity from the wave slope data. The wave slope record was converted to wave amplitude by simply dividing by the wavenumber (also constant), since κA represents the wave slope. A multiplication by the radial frequency ω gave the wave velocity. The analysis of the velocity record was conducted in a manner similar to that for the wave frequency. Each velocity record was segmented into bins containing ten wave periods. The data in each bin then was curve-fitted with a sinusoid to estimate the local wave velocity. While the other wave properties were found to

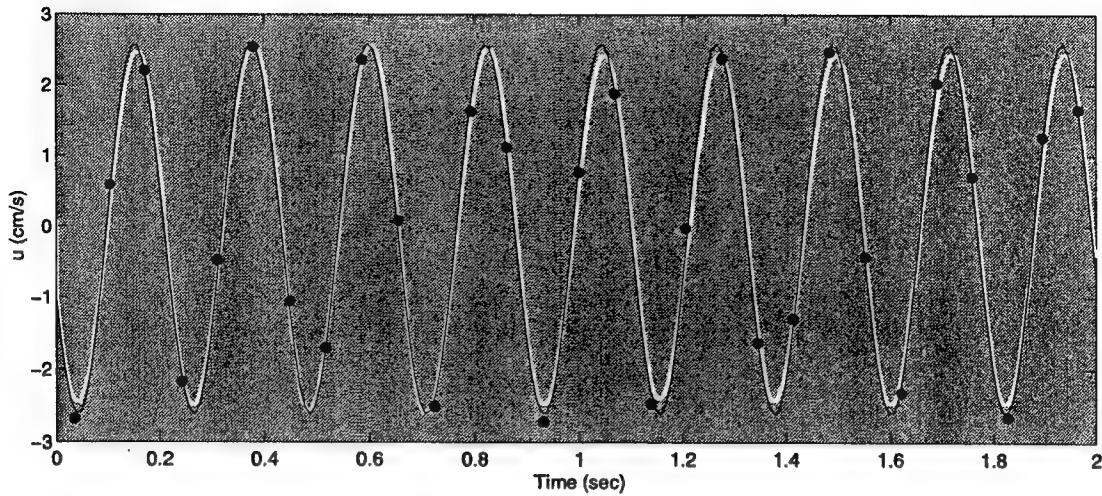


Figure 5-4: Sample wave surface velocity $u(t)$ records computed from the wave slope measurement and the surface PIV velocity measurement. Results are from run 1-1. The white line corresponds to the wave velocity computed using the wave slope data. The black markers correspond to the wave velocity as measured with the surface PIV. The black line corresponds to a sinusoidal curve-fit to the PIV data points. Since the two measurements began at slightly different times, the slope result has been shifted in time to achieve phase alignment of the two records.

be constant over the sampling period, the wave velocity (*viz.*, amplitude) was found to vary somewhat. The PIV velocity record $u(t)$ was binned and curve-fitted as described above for the wave slope data. Figure 5-4 shows an example of the two data sets over a 2 second period.

Because the wave surface velocity varied slightly with time, mean results from each record were less useful for comparison. Since the two data sets were synchronized to within approximately 1/2 second, the binned results were compared point by point. A mean error was defined as

$$\text{error} = \frac{1}{N} \sum_{i=1}^N \frac{|u_{s_i} - u_{v_i}|}{u_{s_i}}, \quad (5.5)$$

where the index i specifies the individual bins and N is the total number of coincident bins. The final results are tabulated in table 5.2. The mean errors in the wave surface velocity were found to range from about 1% to about 7%. In half of the cases, the PIV results were lower than the results from the wave slope. In two cases, the results showed no obvious

Run	Error (%)
1-1	2.09
1-2	5.07
2-1	1.05
2-2	7.15
2-3	2.40
2-4	4.38

Table 5.2: Wave velocity results. The error is the averaged error of the binned results from the wave slope measurements and the surface PIV measurements (equation 5.5).

biases, and in one case, the PIV results were higher.

In the analysis of the wave surface velocity, a measure of the wave amplitude A naturally emerged. The amplitudes of the six runs examined ranged from about 0.1–0.5 cm, and yielded ratios A/λ in the range of 0.01 to 0.04. Based on these values, it was concluded that the use of linear theory was valid.

5.3.4 Sources of error

It was stated earlier that any discrepancies between the results found using the wave slope data and the PIV results would be considered to be errors in the PIV measurement. This was done largely as a convenience in terminology. However, both measurements have uncertainties, and discrepancies between the two measurements could be partially due to errors in the wave slope results. In this final section, some of the possible sources of error for both measurements are discussed.

Comparing the wave frequency results first, these were in excellent agreement, likely indicating that both systems accurately measured this characteristic of the waves. This also implies that the PIV camera frequency is well-known, a useful result itself.

Looking at the wavelength measurement, a small error, about 3%, was detected for the single run examined. Sources of error for this result exclusive to the PIV measurement could be the length conversion used to scale the camera images, PIV processing error, and/or potential shortcomings in the curve-fitting procedure. The effect of the two-dimensional projection of the surface (discussed below) should not play a role in the wavelength estimate.

Chapter 4 looked at possible errors incurred in determining the length conversion and found these to be quite small—tenths of a percent. PIV processing errors are also expected to be small (on the order of 1%), based on the results from chapters 3 and 4. Finally, it is possible there is error attributable to the curve-fitting method of extracting the wavelength.

Since the wavelength estimate using the wave slope data relied solely upon the measured wave frequency and a dispersion relation, the errors incurred should derive from these two aspects of the measurement. As mentioned above, it is believed that both systems accurately measured the wave frequency, therefore errors in the frequency measurement will not be considered. This leaves the dispersion relation that was used to infer wavelength from wave frequency. Issues of finite depth and linear wave theory have already been explored, and should not be sources of appreciable errors. The effect of surfactant rheology on the wavelength–frequency relationship was investigated using the complete short wave dispersion relation derived by Lucassen-Reynders and Lucassen [90]. This dispersion relation includes viscosity, surface tension, and a dynamic surface viscoelastic modulus. Using the measured surface tension, and exploring a range of real-valued elasticities of 0–60 mN/m, a peak wavelength at the run 1-1 frequency was found: $\lambda_{max} = 8.12$ cm. This extremum value still falls short of the PIV result; therefore, surfactant effects are unlikely to be a significant source of the error. Other possible effects that may have slightly altered the relationship between frequency and wavelength could include the fact that there did exist some two-dimensionality in the wave field (*e.g.*, see figure 5-3a), and the influence of the channel side walls. Since no one source of error appears dominant, the measured 3% error is likely some combination of all the aforementioned sources. Physically, the error is less than 3 mm, which can be considered formidable nonetheless.

The errors in the velocity measurements reported in table 5.2 are on the order of 4%. In this case, errors could be due to both the PIV measurement and/or the actual wave slope measurement. Examining PIV first, a possible source of error could be the effect of projecting the three-dimensional wavy surface onto a two-dimensional plane. The data show that the inclination angle of the free surface (θ_1 in figure 5-2) was always less than 0.25 radians. Since the projection due to the PIV measurement involves a multiplicative $\cos \theta_1$ factor, the maximum error introduced would be about 3%. Thus, for these waves,

the effect of the projection should be small, but may explain why the majority of PIV velocity results were near or below the results found using the wave slope data. Additional sources of PIV error are those already mentioned: length conversion and PIV processing errors. Sources of error in the wave slope measurement include the photodetector calibration (output voltage \rightarrow offset position) and the measured distance between the free surface and the photodetector. The photodetector calibration used to determine the laser beam offset yielded 95% confidence intervals of ± 0.016 mm. The calibration therefore is not a likely source of error. The distance between the water surface and the photodetector was measured manually with a scale. This measurement is believed to be fairly accurate, with an estimated error of no more than 1 mm. This uncertainty translates into less than 0.5% error in the estimate of the amplitude. The fact that the distance between the detector and the free surface varies over the passage of a wave was not accounted for in this analysis. For a wave of amplitude 4 mm, the maximum error incurred by assuming a constant distance would be about 2%. Again, no one source of error appears dominant, and the measured errors are likely some combination of a number of these sources.

5.4 Summary

This set of experiments has shown that the surface PIV velocity measurement technique is adept at measuring the flow at a wavy free surface. Wave characteristics—frequency, wavelength, and surface orbital velocity—were measured using surface PIV and compared to independent estimates of these quantities from wave slope data. The measurement of wave frequency appears extremely accurate, and the wavelength and surface velocity field measurements are estimated to possess an accuracy on the order of 4%.

In practice, there will be a limit to the degree of surface waviness that can be accurately measured with the surface PIV technique. Loosely, this limit is constrained by: (1) the focal distance and depth of field, *i.e.*, the image content needs to be reasonably in focus and has to stay relatively near the static position of the surface in order for the length conversion to be valid, and (2) two-dimensional projection errors, *e.g.*, maximum slopes of 1:2, or 26° , are needed to keep projection errors under 10%. Conservatively, this latter condition translates

to $A:\lambda < 1:12$, or with a more relaxed criterion, $A:\lambda < 1:8$. For $A = 2$ cm, these ratios imply a wave with frequency of roughly 3 Hz.

Chapter 6

Oscillating grid-stirred turbulence

6.1 Introduction

The focus of this thesis is to explore the relationship between free-surface turbulence and air–water gas transfer in the laboratory. Outside the laboratory, surface turbulence in lakes and the oceans is generated by several, often interrelated, mechanisms: wind-driven surface shear, waves and wave breaking, and subsurface current interactions. The turbulence generation is therefore at, or near, the water surface. Alternatively, in rivers and estuaries, much of the turbulence is generated from the shear at the bottom boundary. River and estuarine flows are most often simulated easily by open channel flumes that can be pump- or paddle-driven. Simulation of the lake/ocean environment, on the other hand, typically involves wind (-wave) flumes. A potentially significant drawback to *linear*, wind-wave flumes is that they are often very large and consequently are impossible to maintain chemically clean. *Annular* wind-wave flumes, which can be much smaller and are more suitable for “clean” water studies, are additionally desirable because they represent closed systems. For both types of wave flumes, in addition to the turbulent flow, a wind-driven mean flow will exist, which for certain flow measurement techniques, can present complications. Further complications can arise from bottom influences, fetch dependencies, and rotational effects in annular tanks.

When studying the dynamics of an air–water interface and how these dynamics impact the flux of gases across the interface, the above facilities can prove troublesome. Compli-

cations such as physical size, contamination effects, open-loop systems, and the presence of surface waves can potentially obscure the experimental focus and make measurements arduous. A simplified turbulence generation system is therefore desirable in such a case. In this work, the mechanism selected for the generation of turbulence was an oscillating grid-stirred tank. In such a system, turbulence is created below the air–water interface and is advected upward to the surface. The flow produced by an oscillating grid is the result of interactions between the individual jets and wakes created by the motion of the grid bars. These jets/wakes are aligned with the bars near the grid, but eventually interact and break down into turbulence that is carried away from the grid by the jet motions. In concept, no mean flow exists and the turbulence is considered near-isotropic, and homogeneous in planes parallel to the driving grid. (It turns out, this is not necessarily true.) Grid-stirred tanks belong to a larger class of tanks, sometimes referred to as stirred cells. Other commonly used stirrers include rotating impellers and jets (either single or in an array). A drawback to these two methods of mixing is that the produced turbulence is reportedly non-homogeneous (Asher and Pankow [6]). Furthermore, jets require an external pumping device, which can be problematic in terms of introducing possible contaminants. It was for these reasons that a vertically oscillating grid was selected as the stirring device. In addition to favorable turbulence production and lack of dominant mean flows, a grid-stirred tank can be made gas-tight and of a size that renders “clean” water experiments feasible. These features, and the fact that grid-generated turbulence has been the subject of numerous other studies, make such a tank a useful tool for investigating fundamental hydrodynamic processes involved in air–water mass transport.

In this chapter, the design and construction of the grid tank is discussed. The range of grid forcing conditions achievable are outlined and empirical predictions for the generated turbulence are presented. Issues of averaging time, repeatability, and flow stationarity then are addressed. The remainder of the chapter is devoted to describing the flow in the tank. This involves several experimental efforts to quantify both the overall flow structure and the turbulence in the tank. Bulk and surface flow were considered. Mean flows, turbulent energies, and lengthscales are presented for a number of grid forcing conditions. The goal here was not to conduct a completely comprehensive study of oscillating grid-stirred turbu-

lence. Rather, it was to gain a sufficient understanding of the grid tank flow for the range of forcing conditions studied, and to utilize that knowledge to help formulate the experimental plan for the simultaneous air–water gas transfer and surface PIV measurements of chapter 7.

6.2 Grid tank design and fabrication

In designing the grid tank, several factors were strongly considered: (1) physical size, (2) contamination effects, (3) minimal free-surface disturbance, (4) good optical access to all parts of the tank, especially the water surface, and (5) a gas-tight closed-system. In addressing physical size, it was considered important that the tank be large enough to minimize the effect of the side-walls on the desired measurements. In addition, the properties of the turbulence improve as the distance from the grid increases. To handle the issue of cleanliness, strict guidelines were observed for materials selection, as well as for operational practices. In order to avoid free-surface disturbances and provide unobstructed optical access to the water surface, the grid was driven from below the tank. Finally, a closed system was ensured by implementing a sealed tank lid design. These issues are elaborated on in the following sections.

6.2.1 Tank and grid design

There have been a variety of experimental studies using grid-stirred tanks (*e.g.*, Asher and Pankow [6]; Brumley and Jirka [13]; De Silva and Fernando [26]; Hopfinger and Toly [56]; Thompson and Turner [130]). Dating back to 1955 with the work of Rouse and Dodu [113], oscillating grids have been used as a source of (approximately) zero-mean-shear turbulence for studying turbulent mixing and dispersion in homogeneous, stratified, rotating, and two-phase fluids. Grid tanks have ranged in size from 25 cm square by 46 cm deep (Thompson and Turner [130]) to as large as 3.65 m square by 1 m deep (Milgram [101]). The tank designed here was based on the dimensions of the tank used by Brumley and Jirka (also, Chu and Jirka [18]). This was done to allow for the possibility of close comparisons to their results, which included gas-transfer measurements. It also was felt that such a size was

Specification (units)	Symbol	Value
Horizontal dimensions (cm)	—	45.4 × 45.4
Height (cm)	—	57.2
Water depth (cm)	H	50.8
Water volume (l)	V	104.7
Headspace volume (l)	—	15.7
Grid A, B bar width (cm)	d	1.27
Grid A matrix	—	7 × 7
Grid A mesh spacing (cm)	M	6.35
Grid A solidity	—	35.3%
Grid B matrix	—	5 × 5
Grid B mesh spacing (cm)	M	8.89
Grid B solidity	—	25.4%
Grid A, B depth (cm)	$ z_s $	25.4
Grid stroke (twice amplitude) [†] (cm)	S	3.81–11.43
Oscillation frequency [†] (Hz)	f	0–2.5

Table 6.1: Relevant dimensions and parameters for the grid-stirred tank. [†]This range corresponds to the final grid drive design.

large enough to provide a region of flow that was not directly influenced by side wall effects, while still being small enough such that it could be successfully managed for cleanliness. The internal dimensions of the tank were 45.4 cm square by 57.2 cm deep.

Two stirring grids were designed. The form of the most frequently used grid was modeled closely after that of Brumley and Jirka. This grid (grid A) was a 7 × 7 square grid of solid polycarbonate bars. The cross-section of each bar was 1.27 cm square, and the bars were spaced equally in both directions by 6.35 cm. The second grid (grid B) was a 5 × 5 square array of the same bars, spaced 8.89 cm apart in both directions. The mean vertical position of the grid was located midway between the tank bottom and the free surface. All experiments were performed with a water depth of 50.8 cm. For both grids, the distance from the extreme edge of the grid bar ends to the tank side wall was 4 mm. Table 6.1 summarizes the tank and grid design specifications and a photograph of the grid tank is given in figure 6-1.

The tank was constructed entirely from commercial, optically clear, polycarbonate sheet. The walls were made with 1.27 cm thick material, and the bottom from 2.54 cm thick mate-

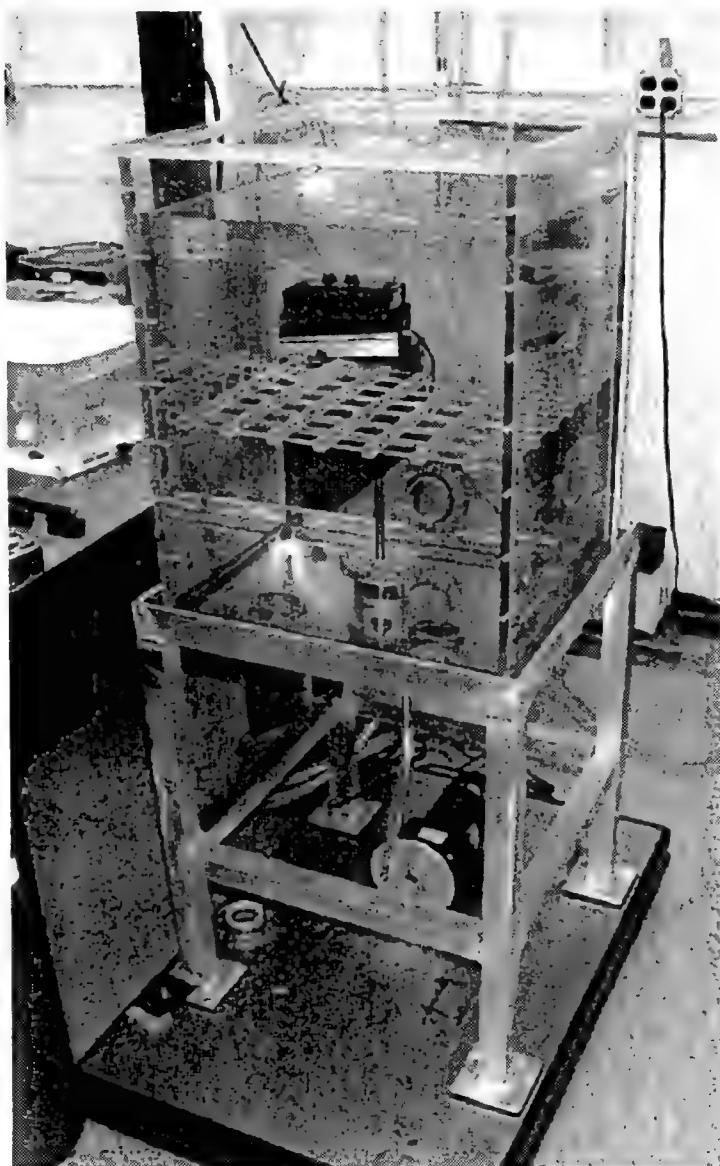


Figure 6-1: Photograph of the oscillating grid-stirred tank in its final form, without lid. The water depth is 50.8 cm, grid A is in place, and the motor-driven grid drive mechanism is pictured.

rial. Polycarbonate was chosen for its excellent chemical resistance and ease of machining. Glass, while being exceptionally inert and ideal for optical access, was not used because of difficulty in machining and bonding issues (see below). All joints were sealed chemically using methylene chloride, leaving no residual solvents. This technique yields very strong joints and does not impart spurious contaminants. A bonding agent like silicon, popular for joining glass, was avoided because of possible leaching effects of the joint seals over time. For additional strength, all of the wall corner joints were reinforced with a vertical array of fastening screws. A mitered rim was installed at the tank top to accommodate a lid and seal. The seal was a custom-made neoprene gasket that mated to the tank rim. The tank lid was made from 1.27 cm thick polycarbonate and had several gas ports installed. These ports allowed for headspace flushing and sampling. A 2.86 cm diameter through-hole was made in the lid to allow for surface tension measurements using the Wilhelmy plate technique described in section 2.3.1. The tank was made gas-tight by locking the lid down on the tank using 14 C-clamps around the rim perimeter.

The tank was designed with the idea of installing a number of measurement probes in the flow. To allow for this, four access ports were created. Two are located in the tank floor and the other two are located on adjacent side walls, above and below the mean grid position. Specifically designed polycarbonate plugs, which can accommodate water-side sensors, mate to the access ports and seal using Buna-N (good chemical resistance) O-rings. Since the ports are all identical and interchangeable, there is flexibility in sensor placement. Sensors can also be placed in the headspace with simple modification to the tank lid as necessary.

The tank was filled with water through the open top. Water could be removed from the tank in two ways. A main drain valve was located in one corner of the tank floor. From this valve, a drain line connected to a pump, which could empty the filled tank in about one minute. Additionally, a standpipe was located in another corner of the tank floor. This standpipe entered the tank through an O-ring seal and was vertically adjustable and removable. The pipe often was used to set the water level, after which it would be sealed with a neoprene (good chemical resistance) stopper and lowered so as to only protrude into the tank by approximately 3 cm. The tank was mounted in a custom-built aluminum frame.

The tank sat in the frame along its bottom edges, leaving the tank underside unobstructed for sensor positioning. The frame was mounted securely to a leveled optical breadboard table that was situated on the laboratory floor. Neoprene pads under the frame feet helped reduce the impact of room vibrations.

6.2.2 Grid drive design

As mentioned earlier, a design goal for the tank was to minimize any free-surface disturbances. In other words, the water surface was to be left free of surface-piercing components. This eliminated the possibility of a grid drive system that was located above the tank as has been employed by many others (*e.g.*, Asher and Pankow [6], Brumley and Jirka [13], De Silva and Fernanado [122]). More importantly, such an arrangement would have interfered with the surface PIV measurement system. This led to the design of a grid drive system located below the tank. For simplicity, a single drive shaft was used to oscillate the grid vertically. This presented the issue of maintaining a watertight seal. In the first design of the tank, a neoprene bellows was used. The grid was fastened at its center vertex to a delrin shaft. The bellows sealed around this shaft and was located in a flooded center well that passed through the tank bottom. The bellows was placed in the well to isolate the flow due to the bellows motion. The well had a fixed top cap with a small through-hole that allowed the shaft to pass through. The other end of the bellows sealed to the removable bottom cap of the well. The well was polycarbonate and had an inner diameter of 7.62 cm. The lower end of the delrin shaft was coupled to a precision-ground stainless steel shaft that ran through a linear recirculating-ball bearing mounted to the underside of the well bottom cap. The lower end of the stainless shaft was connected to the drive mechanism via a flexible coupling.

In the original tank design, an “electric cylinder” was used as the grid drive mechanism. The electric cylinder was a commercially available, compact, linear actuator that operated using a motor-driven lead screw assembly. This device was selected because the necessary control electronics already existed in-house. The actuator had a maximum stroke of 10.16 cm and could be closed-loop, computer-controlled with basic I/O hardware to follow a prescribed displacement path. A small portion of the results in chapter 7 were collected

with the tank using this drive configuration. However, a number of problems arose with this design. First, it became apparent that the bellows was not sufficiently clean. The bellows actually was intended for an automotive use, and the neoprene contained some amount of anti-ozonant to reduce degradation of the material. It was discovered that some of this additive quickly leached out over time. Second, the operation of the bellows caused the tank water level to oscillate by roughly 1 mm, and this was deemed an undesirable complication. Lastly, the electric cylinder lacked the required power to generate energetic grid forcing conditions. For full stroke motion (≈ 9 cm), the peak oscillation frequency achievable was approximately 1 Hz. Consequently, in order to achieve interesting levels of gas transfer and free-surface turbulence, it was necessary to oscillate the grid near the water surface. Other grid tank studies have indicated that the turbulence generated by an oscillating grid becomes homogenized beyond a distance of $2\text{--}3M$ (De Silva and Fernando [26]). Using the linear actuator, the mean grid depth was only 9.53 cm, or $1.5M$. For all of these reasons, a new design was implemented.

The new grid drive design was based on a reciprocating mechanism. A 1/4 HP, parallel shaft, DC gearmotor was used as the driving source. A fabricated aluminum wheel was mounted to the end of the motor shaft. Specifically located through-holes were put in the wheel face to accommodate a removable stainless steel pin. This pin served as the pivot point for an aluminum link arm that coupled to a new stainless drive shaft. A fabricated aluminum yoke and stainless steel pin coupled the link arm and the drive shaft. Rulon sleeve bearings were used in the two link arm pivot holes. Seven through-holes on the wheel allowed for grid strokes (twice the amplitude) of 3.81, 5.08, 6.35, 7.62, 8.89, 10.16, and 11.43 cm. The gearmotor had a maximum rated speed of 165 RPM, or 2.75 Hz (at 7.7 N·m torque). This represented an almost four-fold increase in possible grid speed over the earlier design. The motor was accompanied by a speed controller that could be operated manually or remotely through an analog interface. The relationship between the analog input voltage and the motor speed was determined experimentally, and a voltage calibrator was interfaced to the speed controller to repeatably set the motor speed at a desired value.

The neoprene bellows was eliminated in the drive redesign. It was replaced with a contact seal at the stainless steel shaft. In order to accommodate this new seal, some

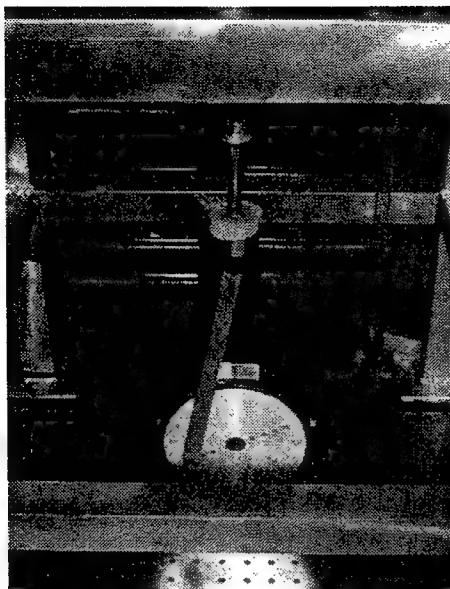


Figure 6-2: Photograph of the final grid drive design. The lower bearing is just visible near the top of the photograph. The grid stroke is set at 10.16 cm

modifications were made to the center well. In final form, the stainless drive shaft was put through two Frelon-lined linear bearings, vertically spaced 20 cm apart. (Frelon is a Teflon-filled composite and is chemically inert.) The lower bearing was mounted to the underside of the well bottom cap, while the upper bearing mounted to the underside of a new, removable upper well cap. The bearing pair ensured that the shaft moved vertically. The upper well cap was designed to house the sealing mechanism. The seal was achieved using a stack of Teflon V-rings. Teflon was selected for its exceptional chemical inertness and low frictional resistance. One, two, or three rings could be installed using appropriate spacers. The rings sat in a recess in the underside of the top cap, and the upper bearing mounting flange was used to put a slight compression on the stack, improving the seal. The Teflon seal was found to work quite well, and it eliminated contamination effects. Additionally, this design reduced the water displacement such that the water level fluctuated by less than $70 \mu\text{m}$. The final version of the drive is shown in figure 6-2. Unless noted, this design was in place for all of the results presented in this thesis.

6.2.3 Grid forcing conditions and empirical turbulence estimates

Oscillating grid-generated turbulence has been the subject of several experimental studies. The works cited thus far represent only a sampling of the numerous references in the literature. The types of experiments, the grid geometries, and the tanks themselves have all been varied. Out of these investigations have come some basic guidelines for grid tank work (De Silva and Fernando [26]). They include: (1) the grid solidity should be less than 40% (to avoid individual jet/wake coalescence), (2) the oscillation frequency should be less than 7 Hz, (3) the end conditions of the grid should be designed properly, and (4) measurements should be made at least 2–3 mesh sizes away. The third guideline requires some explanation. Fernando and De Silva [31] have experimentally shown that a design in which the tank side walls form planes of symmetry with respect to the grid ends yields a flow with reduced secondary mean motions. The authors argue that the attenuation is the result of reducing the Reynolds stress gradients in the fluid. The present final grid tank design meets all of these guidelines.

An often used oscillating grid turbulence reference is the 1976 work of Hopfinger and Toly [56]. Their experimental arrangement consisted of a tank of dimensions 67.5 cm square by 100 cm deep, and grid geometries with mesh sizes of $M = 5$ cm and $M = 10$ cm with $M/d = 5$ in both cases. The grid was located near mid-depth, as this height was found to keep the generation of secondary flows to an acceptable level. Oscillation frequencies of 2–6 Hz and strokes of 4 and 9 cm were studied. Through a series of experiments using a hot-film velocity probe, the spatial decay of the turbulence and the dependence of the turbulence on the grid geometry and forcing were explored. Their results showed the following relationship between the RMS (root-mean-square) horizontal turbulent velocity fluctuation and the grid parameters:

$$u_{HT} = C_1 M^{0.5} S^{1.5} f z^{-1}, \quad (6.1)$$

where C_1 is an empirical constant dependent on grid geometry, and z is the vertical distance from the mean grid position (the mid-plane of the grid at zero displacement)¹. The vertical

¹In equation 6.1, the “virtual origin” for the coordinate z is actually slightly below the mean grid position (Hopfinger and Toly [56]).

velocity fluctuation was observed to be

$$w_{HT} = C_2 u_{HT} \quad (6.2)$$

at sufficient distance from the grid. Furthermore, the horizontally measured longitudinal integral lengthscale was found to increase linearly with distance from the grid:

$$L_{HT} = C_3 z. \quad (6.3)$$

The data of Hopfinger and Toly were best described with the following empirical constants: $C_1 \approx 0.25$, $C_2 \approx 1.1\text{--}1.3$, and $C_3 \approx 0.15\text{--}0.35$. Others (e.g., Brumley and Jirka [13]) have used a value of 0.1 for C_3 . In general, these constants can depend on both the grid geometry and on the stroke. Using these estimates, a turbulent Reynolds number can be defined,

$$Re_{HT} = \frac{2u_{HT}L_{HT}}{\nu}. \quad (6.4)$$

The factor of two arises from the relation for isotropic turbulence, $\ell = 2L$, where ℓ is the standard turbulent lengthscale of the largest eddies, and L is the longitudinal integral lengthscale. A Reynolds number based on the Taylor microscale can be defined by

$$Re_\lambda = (15Re_{HT})^{1/2} \quad (6.5)$$

(see appendix A for details). In an effort to explore grid forcing conditions that spanned a range of Reynolds numbers, a group of grid stroke (S) and frequency (f) combinations was devised. These combinations and the relevant parameters are summarized in table 6.2. Note, that with equation 6.1 and equation 6.3, the Reynolds numbers appearing in table 6.2 are independent of the coordinate z . In the coming sections and chapters, each grid forcing condition will be referred to by its Hopfinger–Toly Reynolds number as given in table 6.2. However, the HT subscript will be dropped, as will the A/B subscript. The latter distinction will be made clear in the text.

S (cm)	f (Hz)	Re_{HT_A}	Re_{HT_B}	Re_{λ_A}	Re_{λ_B}
8.89	0.70	234	277	59	64
6.35	1.40	282	334	65	71
6.35	1.73	349	413	72	79
8.89	1.15	384	454	76	83
7.62	1.55	411	486	79	85
10.16	1.15	469	555	84	91
10.16	1.30	530	628	89	97
7.62	2.25	596	706	95	103
10.16	1.47	600	710	95	103
10.16	1.63	665	787	100	109
11.43	1.50	730	864	105	114
8.89	2.25	751	889	106	115
11.43	1.62	789	933	109	118
10.16	2.20	898	1062	116	126
11.43	2.00	974	1152	121	131

Table 6.2: Grid conditions and characteristic parameters. See text for complete explanation. The constant C_3 was taken to be 0.1. The subscript A denotes grid A, the subscript B denotes grid B.

6.2.4 Tank cleaning and preparation

In this section, the procedures followed for cleaning and preparing the grid tank for operation are briefly described. Some useful guidelines in this area are given by Scott [118]. In all “clean” water experiments, the water used was commercially available spring water. (About 4-1/2 six gallon spring water containers filled the tank to the proper level.) The grid drive components were typically the first to be cleaned. This included the grid, the drive shaft, the well top cap, and the Teflon V-rings. New rings were installed as necessary. Over use, the rings were found to cold-flow and the shaft seal would slowly begin to fail. Typically, two weeks of regular use was the lifetime of a stack of V-rings. During cleaning, all parts were handled with powder-free nitrile gloves. Parts were cleaned with a 50% water, 50% methanol solution and brush. After the 50:50 methanol/water solution cleansing, parts were given a final rinse with the spring water.

The tank itself was cleaned with the grid removed. The rest of the (cleaned) grid drive components would be in place. The first step in cleaning the tank for use was to begin

flushing with laboratory tap water. Both the main drain and the standpipe, in its most lowered position, were used as outlets during this phase. Concurrently, glass cleaner was used on the inside walls and floor with a soft sponge to break down any oily residues. Once the sponging was complete, the tap water flushing would continue for roughly 10–15 minutes. After the tank had been flushed sufficiently to remove any residual cleaning solution, a final rinse and flush with the spring water were performed. Approximately 50 liters of this water were used for the final flush. At this point, the main drain valve would be closed, and the standpipe raised to the proper height. The tank would be filled with the spring water until it reached the standpipe level and the surface water began to drain. While this draining probably did not remove a significant amount of surface-active material, it was done as an additional cautionary measure. The standpipe would be capped with a neoprene stopper (cleaned with the 50:50 methanol/water solution and handled with nitrile gloves) and lowered completely.

In order to remove adventitious surfactant material that always existed in the tank (because no working tank can be 100% free of surface-active substances), a surface aspiration system was employed. A small glass pipette was connected to a tygon tube line that was set up at the laboratory sink to create suction in the line. The pipette was moved by hand around the water surface, drawing up the very upper layer of the interface. This “vacuuming” technique is the same as used by Asher and Pankow [6] in their grid tank experiments. Roughly 5–10 minutes of surface aspiration typically was performed. Airborne contamination was avoided by covering the tank with its lid. Other preparations that involved gas-transfer measurements will be discussed in section 7.5.3.

6.3 ADV measurements of bulk flow characteristics

While the expressions of Hopfinger and Toly are useful for ascertaining a quantitative sense of the turbulence generated by an oscillating grid, variations in tank and grid geometries will result in certain deviations from such empirical formulae. Early in the work with the newly redesigned grid tank, measurements of the tank flow were made using an acoustic Doppler velocimeter (ADV). The goal of these experiments was to quantify the actual turbulence

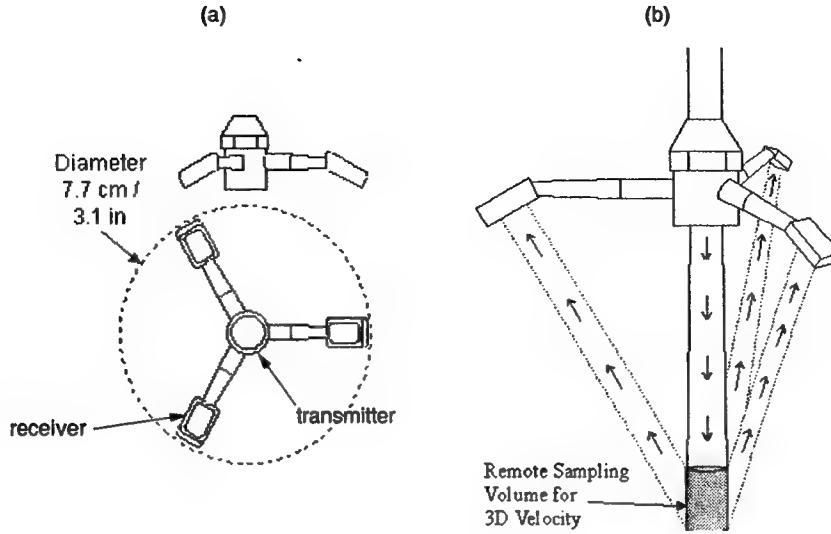


Figure 6-3: Sketch of the ADV sensor probe head. (a) Probe head arrangement showing the center 16 MHz transmitter and three receivers. (b) Sample volume location and size.

levels and gain insight into the degree of secondary mean motions in the tank.

6.3.1 The ADV sensor

The ADV used was a commercially available device similar to the sensor evaluated by Voulgaris and Trowbridge [135]. A thorough description of the sensor principles of operation can be found in that reference. In the present case, the ADV was a 16 MHz device, capable of making remotely-sensed, single-point, three-dimensional velocity measurements in water. The device works on the Doppler shift effect and is configured as a bistatic acoustic Doppler system with a single transmitter and three receivers (see figure 6-3). The three receivers are arranged 120° apart in a circle 7.7 cm in diameter and focus on the sample volume located nominally 5 cm from the transmitter. The sample volume is a cylinder with diameter 4.5 mm and height 5.6 mm. A velocity resolution on the order of 1 mm/s is specified by the manufacturer. The velocity range of the sensor is programmable from ± 3 to ± 250 cm/s, and measurement error is specified as $\pm 1\%$ the measured velocity range. Data can be acquired at sampling rates up to 25 Hz.

For all of the ADV measurements discussed here, the tank lid was removed and the

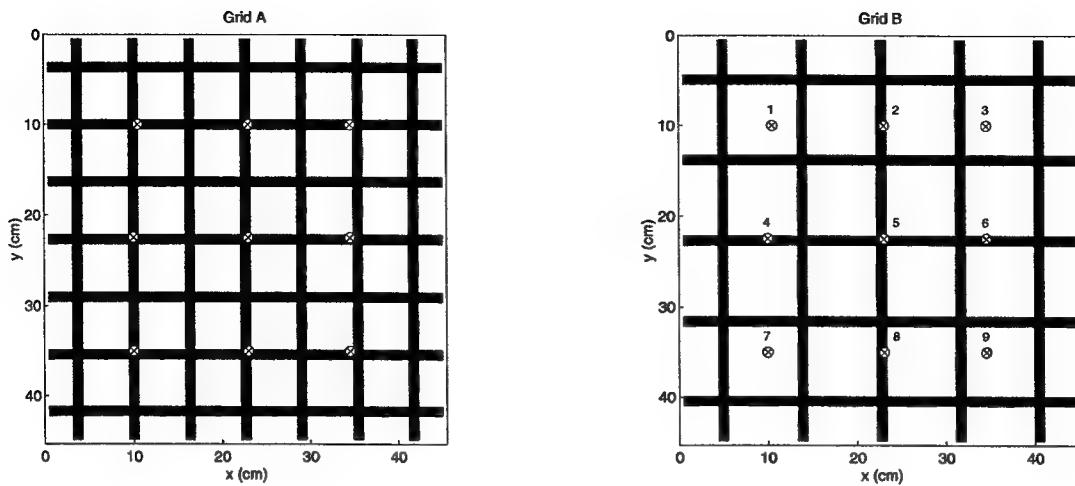


Figure 6-4: Locations of the ADV measurements relative to both grid geometries (overhead plan views). The plot boundaries reflect the tank inner walls. The measurement locations are indicated by numbers 1–9, shown in grid B sketch. Identical numbering applies to grid A. The depth was 8 cm for all measurements. The coordinate system adopted for the grid tank is also shown; the vertical (in/out of page) dimension is measured by z with $z = 0$ at the free surface and positive downward.

sensor was mounted above the tank with the probe head immersed in the water pointing downward. Seeding particles (10–20 μm diameter hollow glass spheres, specific gravity near 1.0) supplied by the sensor manufacturer were introduced into the tank and mixed using the grid motion for a brief period of time. Since the focus of these experiments was bulk flow dynamics, the fact that the sensor pierced the surface was not considered detrimental. Attention to surface cleanliness was also relaxed, as surfactant influences should not penetrate to significant bulk depths. Nine locations in the tank were explored, all at the same depth. The probe head was positioned such that the measurement volume was 8 cm below the water surface, or 17.4 cm above the mean grid position ($2.7M$ for grid A, $2M$ for grid B). Figure 6-4 shows the nine measurement locations in relation to the tank and grid layouts. This figure also introduces the coordinate system used throughout all of the grid tank experiments.

6.3.2 ADV sampling error estimate

In order to determine the ADV sampling error, an experiment was conducted with the sensor in a still tank. This sampling error derives from the inability of the sensor to resolve the Doppler shift and from electronically induced noise (Voulgaris and Trowbridge [135]). As with all of the ADV measurements reported, the data rate was set at 25 Hz. A velocity range setting of ± 30 cm/s was used, as this was the most common setting employed for the ADV measurements in this thesis. Two measurements were made. For the first (experiment A), the sensor was set in the tank in an arbitrary orientation and position, and the tank was seeded and left overnight, with the measurement being made the following morning. The second measurement (experiment B) immediately followed the first. In this case, the ADV was moved delicately from its arbitrary location and mounted in position 1 at the 8 cm depth. This relocation was done slowly so as to induce minimal flow motion. Fifteen minutes after the relocation was completed, the measurement for experiment B was made. In both experiments, a 15 minute time series at 25 Hz was taken (see section 6.3.3 for length of time series information). Table 6.3 summarizes the results of both experiments. The standard deviations of the velocity components were calculated and are reported in table 6.3. The errors measured during experiment B were roughly half that of experiment A. Correspondingly, both the signal-to-noise ratios (SNRs), and the correlations (indicative of measurement quality) are much improved in experiment B over that of A. The measurement of the velocity component normal to the transmitter axis, w , shows the smallest error compared to the other two components. The difference is close to the approximate factor of five that is to be expected for this bistatic tip geometry (Voulgaris and Trowbridge [135]). The total velocity error estimated from experiment B falls within the specified value from the manufacturer, 0.30 cm/s at this velocity range.

6.3.3 Velocity time-record requirements

In this section, experiments that explored the characteristics of the ADV sampling record necessary to determine stable temporal flow statistics are described. The first requirement determined was the length of time needed for the grid tank flow to come to an equilibrium state from a rest state. Fifteen minutes was assumed a sufficient warm-up period, and

Experiment	Vel. range (cm/s)	SNR (dB)			Correlation (%)			Std. dev. (cm/s)		
		<i>u</i>	<i>v</i>	<i>w</i>	<i>u</i>	<i>v</i>	<i>w</i>	<i>u</i>	<i>v</i>	<i>w</i>
A	± 30	14	13	6	96	96	89	0.59	0.82	0.14
B	± 30	22	22	13	99	99	97	0.34	0.43	0.08

Table 6.3: Results from calm water ADV measurements. SNR refers to the characteristic values noted during the experiment from the ADV acquisition software. Correlation refers to the correlation coefficient values for each velocity component; a correlation value greater than 70% is required for accurate measurements.

measurements indicated that this was indeed sufficient. One run with grid B at $Re = 555$, and two runs with grid A at $Re = 469$ and 600 (grid frequency setting was not correct) were examined. Measurements were made at location 1. Data from each run were collected from the onset of grid motion, and at least 45 minutes of data were logged. Since the three data sets yielded the same conclusion, only one is discussed. Figure 6-5 shows the temporal behavior of the flow over a 45 minute period for $Re = 469$. Local time-averages of each component of the total velocity and local RMS values of the velocity fluctuations were calculated over the full record. The local bin length was set at 200 grid cycles, or just under 3 minutes for this grid forcing. This time was felt to be sufficiently long to contain a large number of integral scales. (Integral time scales for the grid turbulence were found to be on the order of one second). It appears all components of the flow come to an equilibrium state shortly after the motion begins. There is no evidence of a lengthy start-up period. In fact, given the observed variability over the measurement period, it would seem that the flow becomes fully developed within the first few minutes. While such a result may not strictly hold for all grid forcing conditions, the 15 minute warm-up period that was observed for all ADV experimental runs should have been more than sufficient.

The second requirement determined was the length of the ADV record (after an equilibrium state had been reached) necessary to calculate meaningful estimates of statistical parameters. Again, 15 minutes was taken as a sufficient length of time, and this was verified through a number of experiments. It is helpful to consider some statistically derived estimates that can be obtained for both mean values and RMS turbulent values. Tennekes and Lumley [129] have shown that the mean square error in the estimate of a mean value

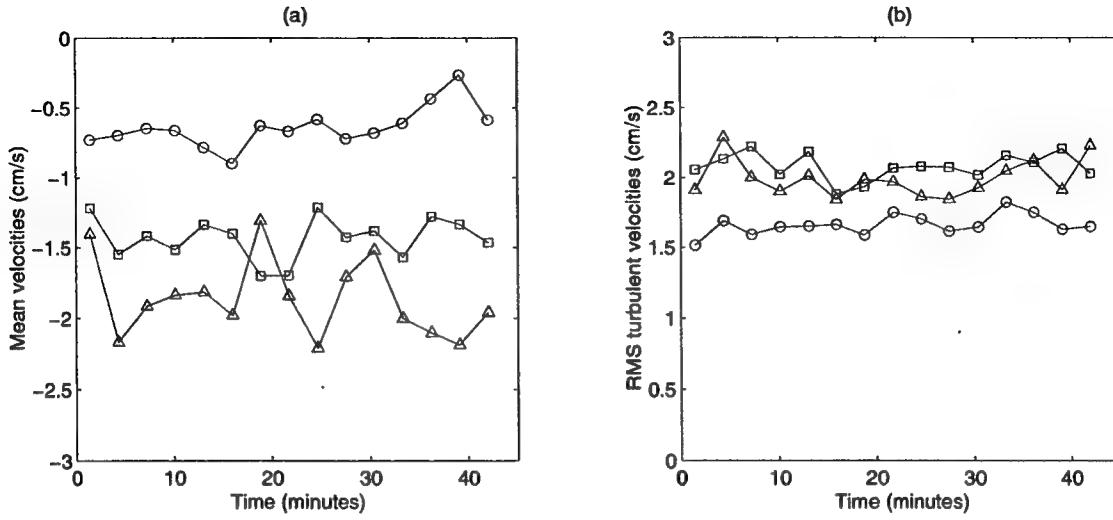


Figure 6-5: Temporal behavior of flow statistics from the onset of grid motion measured at location 1 for $Re = 469$. The grid oscillation started at time equal to zero. (a) Mean flow estimates, and (b) RMS turbulent fluctuations. Squares correspond to u components, circles correspond to v components, and triangles correspond to w components. Each point represents a local time statistic over 200 grid cycles, or just under 3 minutes.

from a finite record length can be approximated by

$$\overline{(U_{T_r} - U)^2} \cong 2\overline{u^2}(\tau/T_r), \quad (6.6)$$

where U_{T_r} is the mean obtained from the sample record of length T_r , U is the true mean, $\overline{u^2}$ represents the variance of the velocity record, and τ is the integral time scale. (Overbars in this case refer to time averages.) Using long time series (30–90 minutes) for grid A runs at $Re = 282$ (location 7), 469 (location 1), 530 (location 7), and 974 (location 7), and a grid B run at $Re = 555$ (location 1), equation 6.6 was used to provide error estimates for a finite record length. The integral time scale was estimated from the auto-correlation of the velocity fluctuation times series; typical values were 0.5–1.3 seconds. The estimated and measured results for 15 minute record lengths are tabulated in table 6.4. Values estimated from equation 6.6 (taking the square root) are consistent with the measured values. Both results indicate that the RMS errors for the estimates of the mean using a 15 minute time series are on the order of 1 mm/s, which is quite acceptable.

Re	RMS($U_{T_r} - U$) (cm/s)		RMS($V_{T_r} - V$) (cm/s)		RMS($W_{T_r} - W$) (cm/s)	
	estimate	measured	estimate	measured	estimate	measured
282	0.04	0.09	0.06	0.00	0.09	0.02
469	0.07	0.05	0.10	0.03	0.13	0.08
530	0.07	0.07	0.08	0.03	0.14	0.11
555	0.08	0.09	0.06	0.11	0.09	0.13
974	0.08	0.03	0.09	0.06	0.13	0.13

Table 6.4: RMS errors for mean values determined from a finite record length of 15 minutes. The estimated RMS errors were calculated from the square root of equation 6.6. The measured errors were computed by using the means from the first 15 minutes of the time series (after reaching steady state) and assuming that the true means could be considered equal to the full time series means.

Considering the turbulent velocities, Bendat and Piersol [9] provide the following minimum record length (in seconds) required to achieve a specified error in estimating an RMS value of a sampled record $x(t)$ having a true mean value of zero:

$$T_r = \frac{1}{4B\epsilon^2}, \quad (6.7)$$

where B is the bandwidth (in Hz) occupied by $x(t)$, assuming a uniform spectrum within B . The normalized standard error ϵ is defined as

$$\epsilon = \frac{\text{s.d.}[\hat{\Psi}_x]}{\Psi_x}, \quad \text{s.d.} \equiv \text{std. dev.}, \quad (6.8)$$

where $\hat{\Psi}_x$ represents the RMS value estimate, and Ψ_x is the true RMS of $x(t)$. This expression can be used to calculate a required record length for the velocity fluctuations, since their means are identically zero and we are interested in RMS values. For the present case, where the record does not have a uniform spectrum, the estimated required record length will not be strictly correct, but can still serve as a useful check. For a bandwidth of 12.5 Hz, and requiring $\epsilon = 0.01$, $T_r = 3.3$ minutes. Using the data sets that were employed previously to look at the necessary record length for mean flow values, the normalized standard error (equation 6.8) was calculated along each time series in independent, local 3.3 minute bins. The true RMS was taken as the full time series RMS. For all grid forcing

conditions studied, this error was 2–5%. Errors of 1% were not obtained because of the idealized model used in deriving equation 6.8. For a 15 minute time average, the errors should be reduced by a factor of approximately $\sqrt{5} \cong 2.2$, based on equation 6.7. It therefore can be concluded that the choice of a 15 minute record length was more than sufficient.

6.3.4 Flow repeatability

The next diagnostic explored for the grid tank flow was repeatability. Ultimately, for the goals of this thesis work (chapter 7), flow repeatability was not crucial since measurements were made that characterized the desired flow parameters for every experimental run. However, it was desirable to be able to characterize the bulk turbulence levels for the various grid forcing conditions and utilize these estimates for all subsequent experiments without having to repeat such measurements for each instance. This required an appropriate degree of flow repeatability.

A specific experiment was conducted to ascertain an estimate of the repeatability of the grid tank flow. Grid B was used, and $Re = 555$ was studied at location 1. Ten experimental runs were conducted. The procedure was to start the grid oscillating, wait 15 minutes for a steady state to be reached, take 15 minutes of velocity data, and then wait 45 minutes for the fluid motion to dissipate. This procedure was followed consecutively for runs 1–6, and also for runs 7–10. However, between runs 6 and 7, the wait time for the flow to dissipate was approximately 2 hours and 45 minutes. Figure 6-6 shows the results of the ten runs. It is evident that the repeatability of both the mean flow and the turbulence are quite good within each grouping of runs, 1–6, and separately, 7–10; the maximum deviation from the mean of a group is approximately 10% for mean velocities, and only 3% for the RMS fluctuations. However, there is a distinct change in behavior between runs 6 and 7. It appears that the additional wait time of 2 hours induced a somewhat different flow. The most dramatic variability in the mean flow occurred for U , which underwent a change of nearly 1.3 cm/s. The greatest variability in the turbulent components occurred for v , which changed by approximately 0.4 cm/s ($\approx 20\%$). It is believed that this variability is due to a sensitivity of the flow to initial conditions. This type of behavior was also reported by

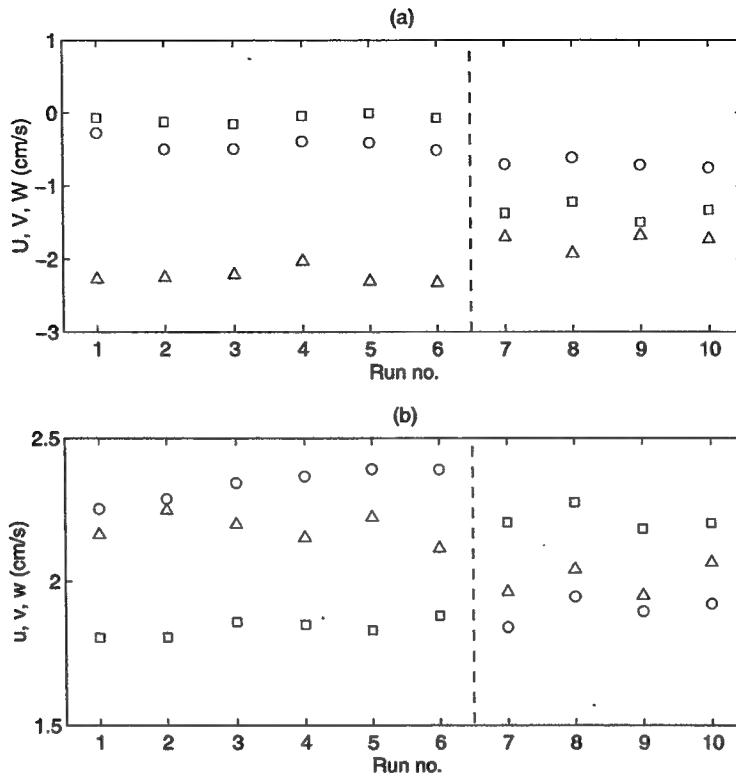


Figure 6-6: Repeatability measurements for $Re = 555$. Grid B was used and all measurements are at location 1. (a) Mean flow estimates, and (b) RMS turbulent fluctuations. Squares correspond to u components, circles correspond to v components, and triangles correspond to w components.

McDougall [93], and is mentioned in passing in the work of E and Hopfinger [30]. In a much smaller grid-stirred tank (25.4 cm square and 46 cm deep), McDougall encountered mean motions that were found to depend on the initial state of the tank. In particular, McDougall describes how the addition of an LDV seed particle mixture gave rise to a large scale mean circulation that “became a permanent feature of the flow.” Furthermore, the location of where the seeding mixture was introduced was found to have an effect on this mean circulation. One result showed that the mean vertical velocity at a particular point in the tank changed by nearly a factor of eight when the seeding mixture was poured in at a different location; the RMS vertical velocity changed by approximately 25%. McDougall also notes that such mean motions were found to persist for very long periods of time. As an

explanation, McDougall points to the nonlinear character of the formation of a mean flow by an oscillating grid. Therefore, slight differences in the ambient flow could give rise to very different mean circulations. To address the results of figure 6-6, the following possible explanation is offered. When the first run was performed, a certain flow pattern was set up, and subsequent trials were “locked in” to that particular flow. Only when the time between the trials was very long could a different flow be set up that did not lock in to the earlier flows.

Another group of results can be assembled together to give further indication of the degree of repeatability. Figure 6-7 shows a collection of ADV measurements for a range of grid forcing conditions where multiple measurements were available so as to give a sense of flow repeatability. Grid A results for $Re = 234, 349, 469, 596, 730$, and 898 are presented. At least three separate measurements were available at each condition. For two of the groups of measurements, only 8 minute records were available. Considering the mean components, the difference between extreme values is always less than 1 cm/s, excluding the group for V at $Re = 898$. For the turbulent fluctuations, the difference between extremes is also always less than 1 cm/s, with the majority of differences being less than 0.5 cm/s. A rough estimate would indicate that, for grid A, the mean values may vary by 25% and the turbulent fluctuations may vary by 15% from any given run to any other given run. This result is slightly better than that found for $Re = 555$ using grid B. It may therefore be possible that one grid is preferable to the other in terms of flow repeatability. In either case, the levels of repeatability were deemed acceptable for the present purposes.

6.3.5 Stationarity of the flow

The final diagnostic investigated for the grid tank flow was stationarity. Since the forcing of the flow is periodic and time-invariant, the bulk flow would be expected to exhibit reasonably stationary behavior. However, the same can not necessarily be said for the flow at the free surface because of time-dependent surfactant effects (*e.g.*, adsorption). Free-surface behavior is not addressed in this chapter, but is discussed in chapter 7.

In order to test for stationarity, the guidelines outlined in Bendat and Piersol [9] were followed. Three lengthy time series were chosen for analysis. All three runs were with

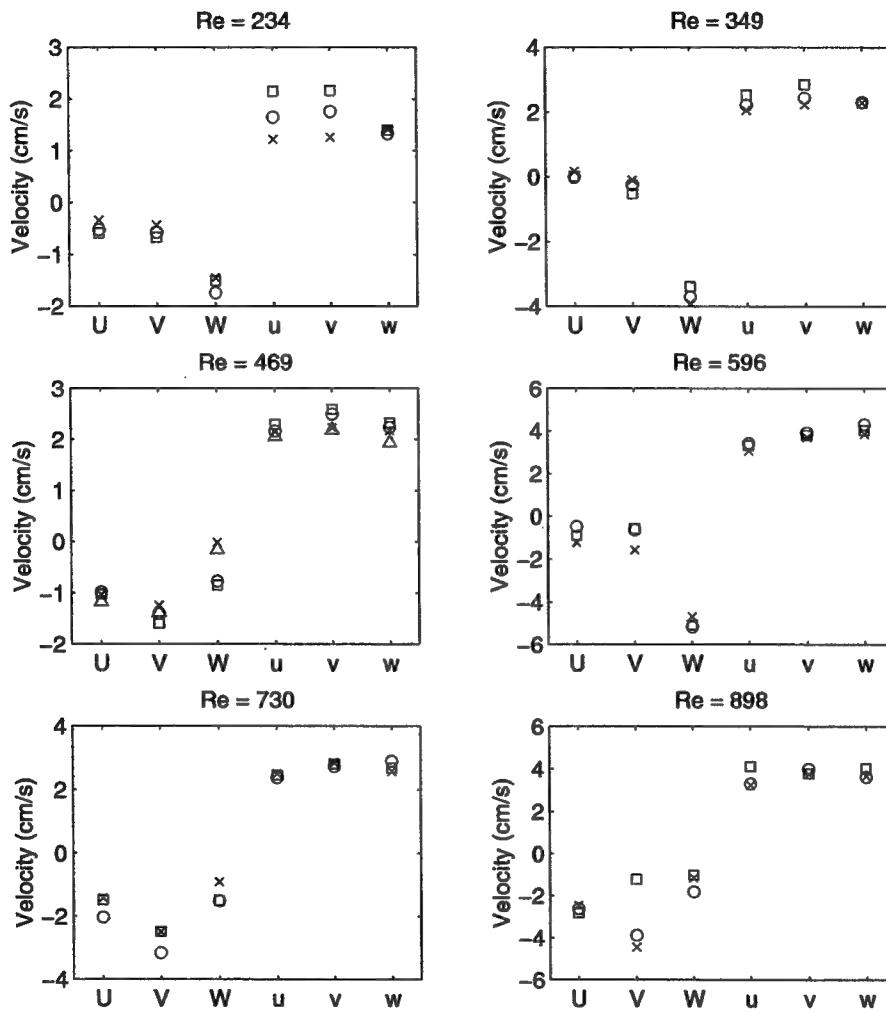


Figure 6-7: Collection of ADV measurements indicating degrees of repeatability. The measurements are for grid A and are at location 1. The circles and crosses are taken from 8 minute records, while the squares and triangles are from 15 minute records. Some measurements were made on the same day, others were separated by several days.

grid A, at position 1, and were for $Re = 282, 530$, and 974 , to provide a wide range of conditions. Each run was at least 90 minutes, post warm-up time. In similar fashion to the approach taken in assessing the length of the development time discussed in section 6.3.3, local time bins were used to compute flow statistics over the entire time record. Each local bin was equivalent to 200 grid cycles as before. The quantities investigated were the mean and mean square of the total velocities, and the RMS and mean square of the turbulent velocity fluctuations. Shown in figure 6-8 are the results for $Re = 974$. The other two grid conditions behaved similarly. These results indicate stationarity of the flow—no underlying trends are obvious, and the variability observed in the flow statistics is within the expected variability due to sample variations. In addition to inspecting the time series qualitatively, two nonparametric tests that do not require special knowledge of the data (*e.g.*, frequency composition) were performed. These tests, described in detail by Bendat and Piersol, are the *run test* and the *trend test*. Both tests are based on a hypothesis that a sequence of measured values (*e.g.*, the mean values and mean square values), are each independent sample values of a random variable with a given mean value and a given mean square value. If the hypothesis is true, the variations in the sequence of sampled values will be random and display no trends. Consequently, the behavior of the sampled values will be as expected for a sequence of independent random observations of the random variable. The run and trend tests are designed to test the above hypothesis. Both tests were used on the sampled data resultant from the locally binned calculations (the data shown in figure 6-8, for example). All of the sequences, with the exception of a negligibly small few, yielded accepted hypotheses of stationarity at a 0.05 level of significance. It therefore was concluded that the oscillating-grid *bulk* flow could be considered stationary.

6.3.6 Spatial variability of the flow

In this section, the primary results from the ADV measurements are presented. As mentioned earlier, the goal of these experiments was to provide a first order estimate of the nature of the grid tank flow. The interest was to quantify the actual turbulence levels, determine the significance of secondary motions, and explore the spatial homogeneity in a surface-parallel plane. Most of the measurements were made for the grid A configuration.

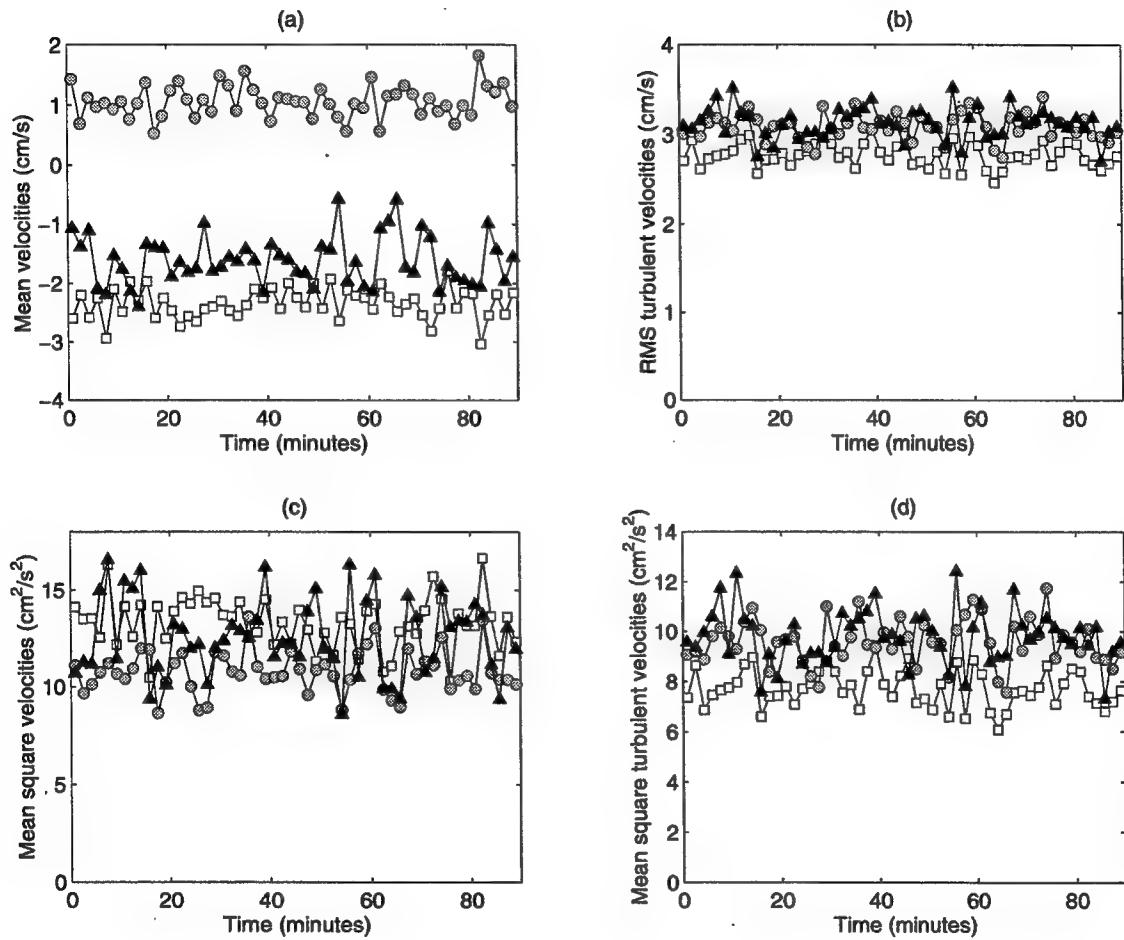


Figure 6-8: Locally binned flow statistics for $Re = 974$ (grid A). The measurements were made at position 1. Squares correspond to u components, circles correspond to v components, and triangles correspond to w components. Each point represents a local bin result over 200 grid cycles.

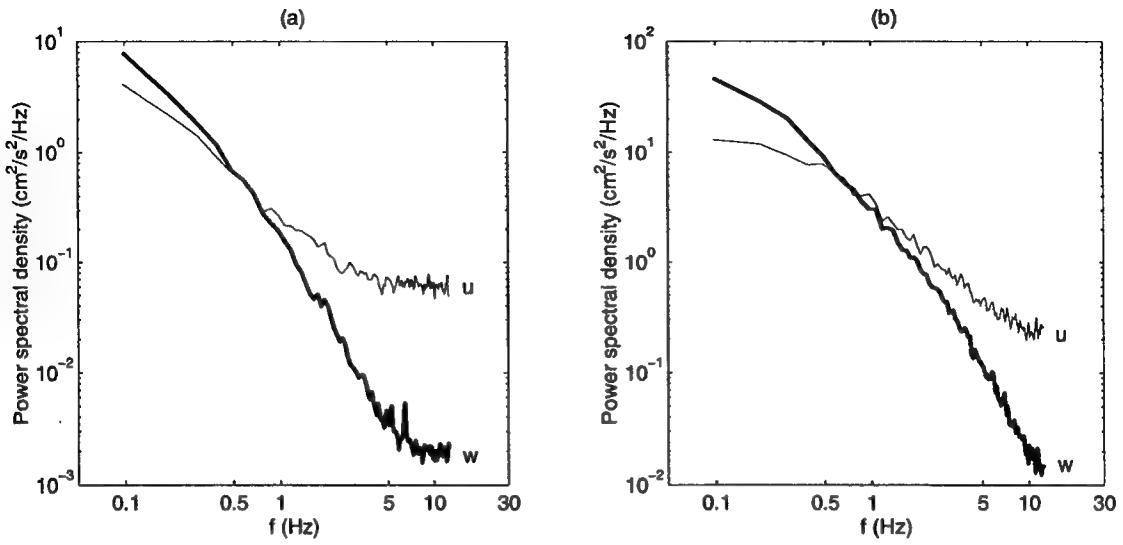


Figure 6-9: Turbulent velocity spectra showing the influence of the ADV noise floor. (a) $Re = 469$, grid A, location 2. The filter cut-off frequency used for u (and v) was 5 Hz; the cut-off frequency for w was 9 Hz. (b) $Re = 596$, grid A, location 9. The data were left unfiltered.

ADV measurements were made at the nine locations shown in figure 6-4, at a depth of 8 cm, for all 12 forcing conditions with grid A, and for $Re = 334, 706$, and 864 with grid B. For these experiments, the grid was started 15 minutes prior to the first measurement to assure a steady state. Fifteen minute velocity time series were collected at each location. After a measurement was completed, the sensor would be moved to the next location, and the next measurement begun. The grid was oscillated continually during all nine measurements. At least 45 minutes separated different grid forcing conditions. Upon examination of the velocity time series, a noise floor for the ADV sensor emerged. Figure 6-9 shows spectra from two runs that reveal the presence of the noise floor. To prevent sensor noise from augmenting estimates of the RMS turbulent fluctuations, the velocity time series were filtered before analysis. A third-order, low-pass Butterworth filter was used to remove the higher frequency noise in the data. Filter cut-off frequencies were estimated by visual inspection of the raw spectra. The two horizontal velocity components were treated collectively since their spectral characteristics were very similar. Therefore, for each run, two filter cut-off frequencies were used, one for u and v , and one for w . When an obvious noise floor was not

evident (e.g., figure 6-9b), the data were left unfiltered. The spectral energy at the noise floor for the horizontal velocity components was found to be approximately 30 times that of the energy at the vertical velocity noise floor. This result is consistent with the sensor design (Voulgaris and Trowbridge [135]). The low pass filtering operation had no effect on the mean velocities (as expected), and reduced the RMS turbulent velocity estimates by roughly 5–40% (horizontal) and 0–5% (vertical), with the largest reduction occurring for the lowest Reynolds numbers.

Another approach to removing the noise influence was explored by looking at the calm water data from section 6.3.2. The spectra of these “no flow” time series showed white noise across the sensor bandwidth. If the noise during the grid flow runs was also uniformly distributed and of the same amplitude, then as a first approximation, the variance of the noise signal during the stagnant condition can be removed from the variance of the turbulent fluctuations during the grid-forced runs. Using experiment B results from table 6.3 yielded RMS turbulence estimates that were in good agreement with those found from filtering. Since it allowed for noise removal on a run-by-run basis, the filtering procedure was believed to be a more robust approach. Therefore, the filtered results were used for the subsequent analysis.

The mean flow results for all of the grid A runs are shown in figures 6-10 and 6-11. These figures reveal the mean flow behavior of the horizontal velocity components and the areas of mean upwelling and downwelling. For forcing conditions with strokes less than 8 cm ($Re = 282, 349, 411$, and 596), the horizontal mean flows are very non-symmetric and unorganized when compared to the flows with larger strokes. These larger stroke results show consistent, outwardly directed mean flows in the horizontal plane at 8 cm depth. It is also found that the mean vertical velocities at the measurement locations are almost always directed toward the free surface. By conservation of mass, this would imply that a significant downward return flow must exist somewhere else in the tank. For the cases where the stroke was greater than 8 cm, the directions of the horizontal mean flows would indicate that the return flow exists at the tank walls. This was verified with simple dye visualization and more rigorously with subsequent PIV measurements (section 6.4). Significant mean upwelling motions often are found in the regions of the tank corners. These upwellings can

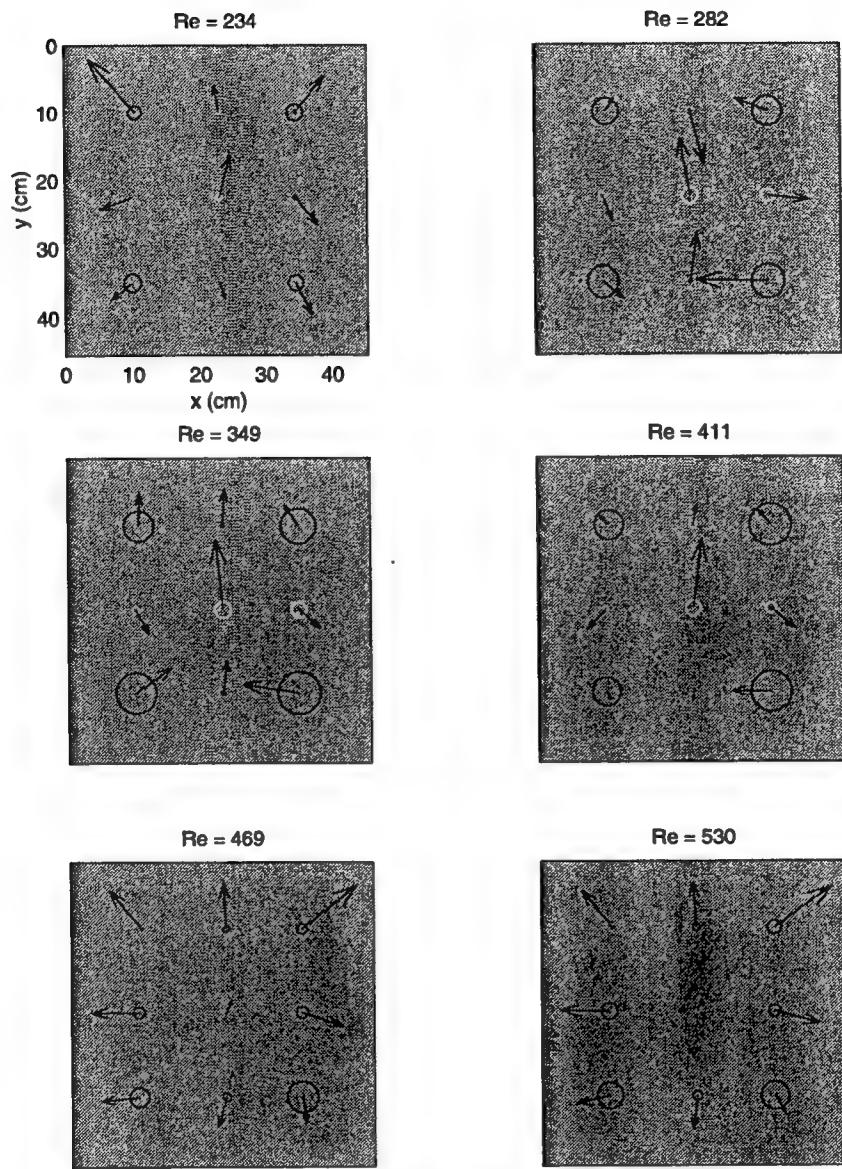


Figure 6-10: Spatial variability of mean flow components measured at a depth of 8 cm for grid A. Each panel domain represents the plan section of the grid tank, 45.4 cm on a side. The vectors show the mean flow in the xy -plane. The circles are indicative of the magnitude and direction of the mean vertical velocity (in/out of page). The circle diameter is proportional to the velocity magnitude; circles in black are upward velocities, circles in white are downward. The circle scaling is consistent for all panels; the vectors are scaled separately for each panel for visibility.

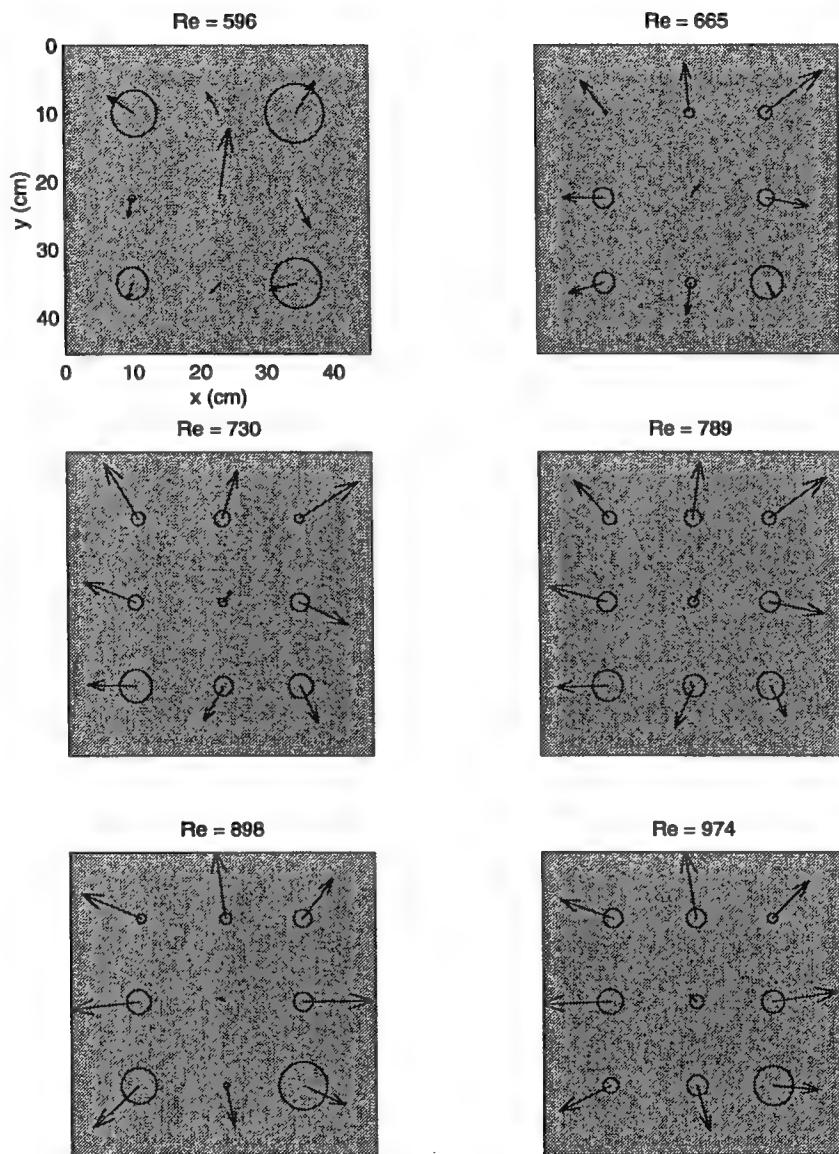


Figure 6-11: Spatial variability of mean flow components measured at a depth of 8 cm for grid A. See figure 6-10 caption for details.

be quite strong as in the case for $Re = 596$. The upwelling regions become less localized for larger Reynolds numbers. The center area of the tank shows the least mean motion in general. In a number of the low Reynolds number cases, the center is a region of mean downwelling. Based on these coarse measurements, a picture of the overall mean flow can be developed. Generally speaking, it would appear that above the grid, a mean flow exists that consists of four upward jets in the tank corner regions. The fluid carried by these jets is returned downward predominantly along the tank walls, and in some cases, is also returned within a downward jet located in the middle of the tank.

The turbulence results for all grid A runs are presented in figures 6-12 and 6-13. Some similarities exist between the spatial distributions of the turbulent velocities and the mean flow velocities. The corner locations show the largest RMS turbulent velocities, while the center shows the smallest. As the Reynolds number becomes higher, this spatial variability is homogenized to a degree. These figures also give some indication of the local isotropy of the turbulence; this is dealt with further below.

The results for the three runs using grid B are shown in figure 6-14. Notable from these results is the appearance of somewhat increased homogeneity over that found for the grid A results. In all three cases, the mean upwelling motions are less localized, and the turbulence shows greater spatial uniformity. This is somewhat surprising since the measurement plane is located fewer grid mesh sizes away than the situation for grid A, and *less* homogeneity would be expected.

To gauge the intensity of the turbulence relative to the mean flow, the ratios u/U , v/V , and w/W were examined. These are shown in figure 6-15. In general, the turbulent velocities are 2-8 times their respective mean velocities. The smallest ratios occur for the vertical velocity, where there is typically a steady mean vertical velocity to yield the reduction. In some cases, zero-mean-flow is approached and the turbulence dominates the velocity measurement, producing ratios that are very large.

Spatially averaged values of the RMS turbulent velocities are shown in figure 6-16 as a function of grid Reynolds number. This figure shows the general increasing trend of the turbulent velocity as the empirically predicted Reynolds number increases. For all grid forcing conditions, there appears strong isotropy in the horizontal velocity components,

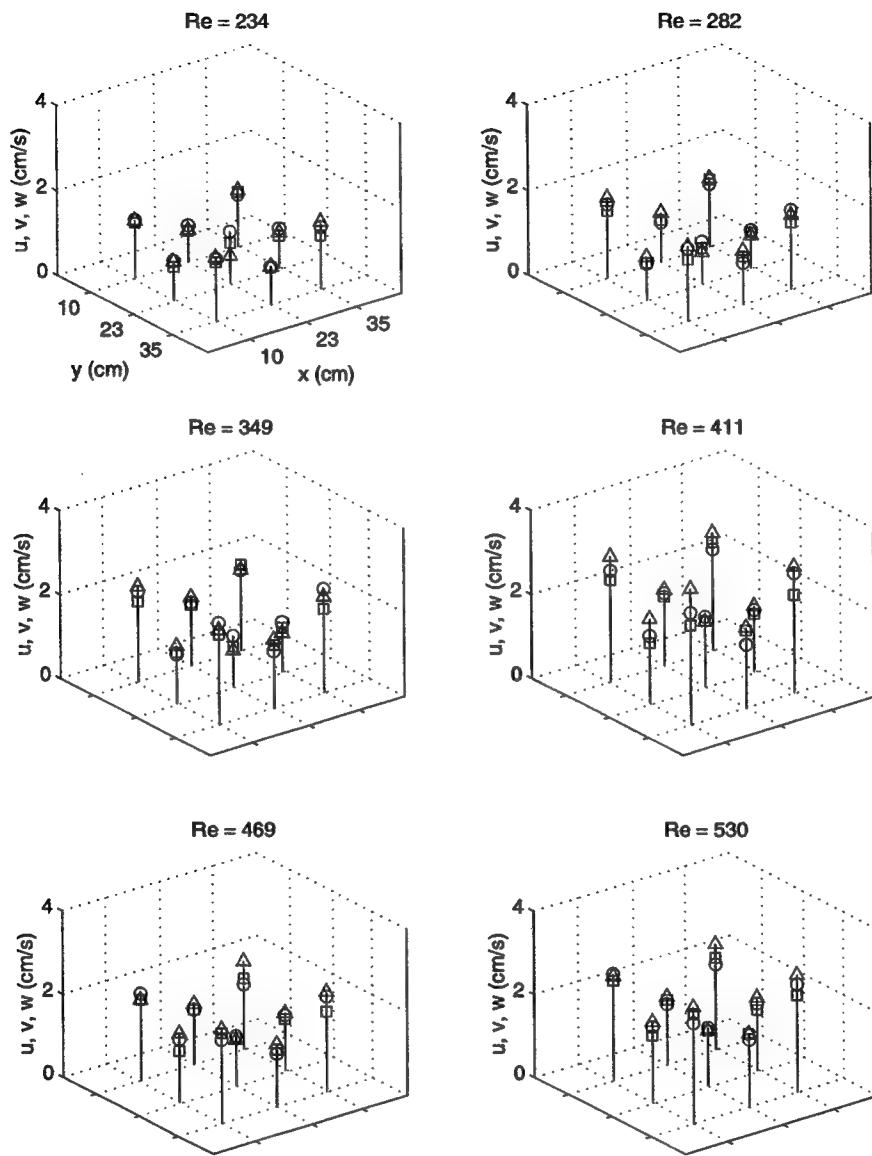


Figure 6-12: Spatial variability of turbulent fluctuating velocities measured at a depth of 8 cm for grid A. The xy -plane is at a depth of 8 cm, and covers the complete plan section of the grid tank. Squares correspond to RMS u velocities, circles correspond to RMS v velocities, and triangles correspond to RMS w velocities.

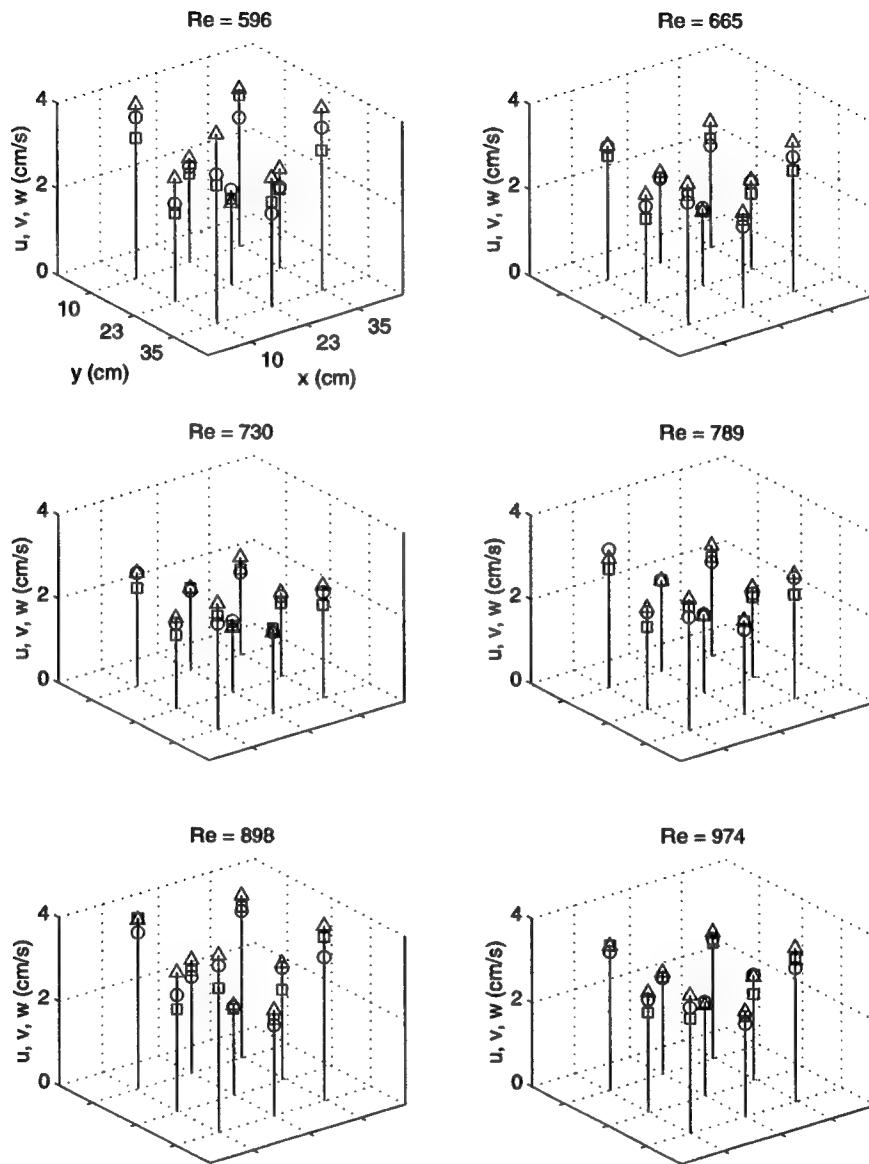


Figure 6-13: Spatial variability of turbulent fluctuating velocities measured at a depth of 8 cm for grid A. See figure 6-12 caption for details.

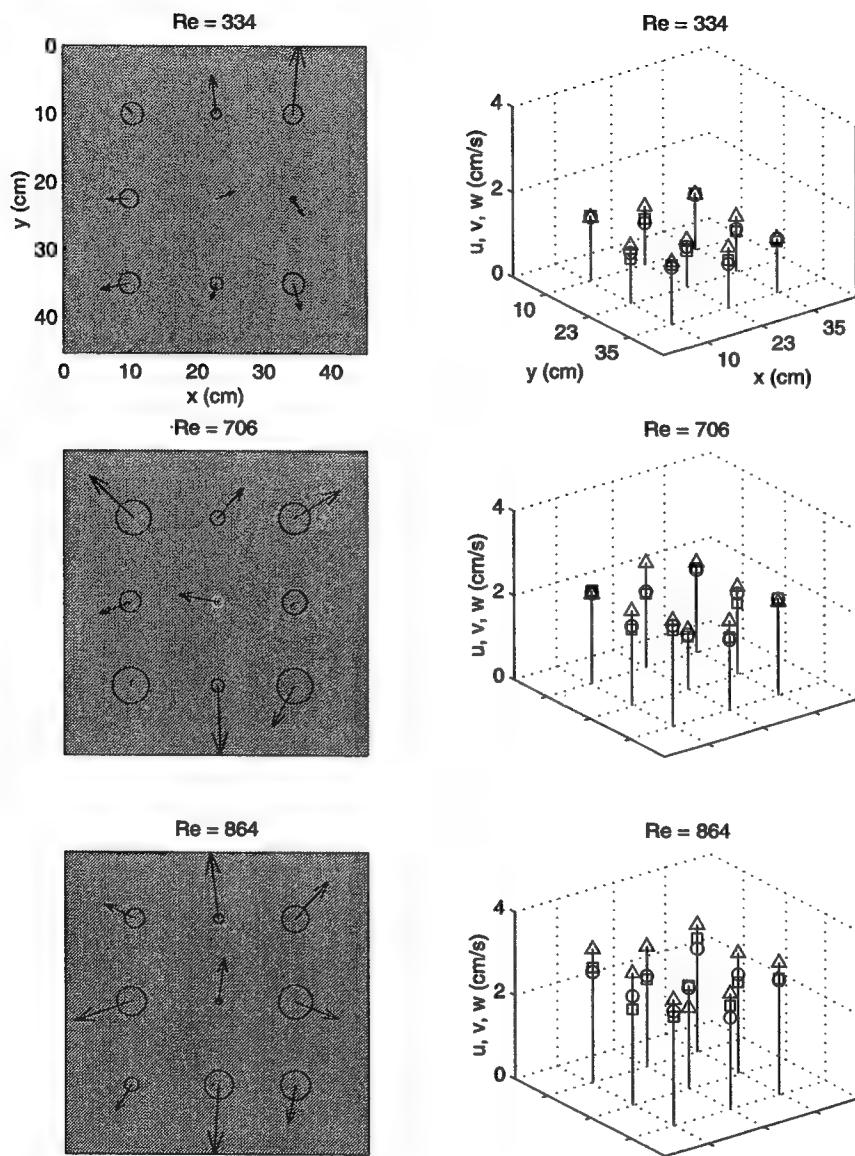


Figure 6-14: Spatial variability of mean and turbulent fluctuating velocity components measured at a depth of 8 cm for grid B. See figures 6-10 and 6-12 for caption details.

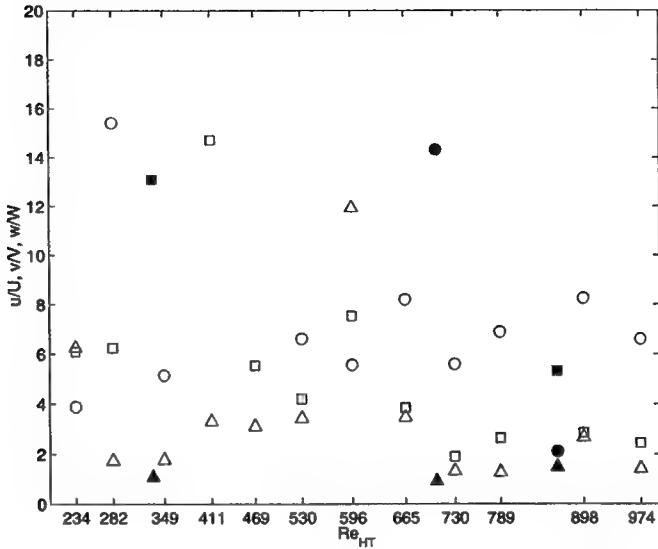


Figure 6-15: Ratios of RMS turbulent velocity fluctuations to mean velocities as a function of grid Reynolds number. The points reflect the spatially averaged values (over the nine measurement locations) of the local velocity component ratios, u/U (squares), v/V (circles), and w/W (triangles). Some ratios were greater than 20 and are not shown here. The solid markers indicate grid B results; all others, grid A.

and the predicted weak anisotropy for the vertical component. To further assess the degree of isotropy of the turbulence, ratios of the spatially averaged turbulent velocities were computed and are shown in figure 6-17. These data indicate a ratio between the RMS vertical turbulent velocity and the RMS horizontal turbulent velocity of approximately $C_2 = 1.1$. This value agrees nicely with what has been found in other studies.

The empirical relations of Hopfinger and Toly were explored next. Shown in figure 6-18 is the relationship between the predicted RMS turbulent velocities and the observed velocities. The predicted velocities were determined from equation 6.1 using $C_1 = 0.25$, the appropriate grid mesh, and the frequency-stroke combinations. All three measured velocity components are included in this figure. In order to include the vertical component, these values were scaled by the estimated C_2 (above) for the present system. The agreement appears slightly dichotomous. Above an RMS velocity of about 2 cm/s, the observed results agree quite well with the empirical predictions. Below 2 cm/s, the observed velocities are consistently

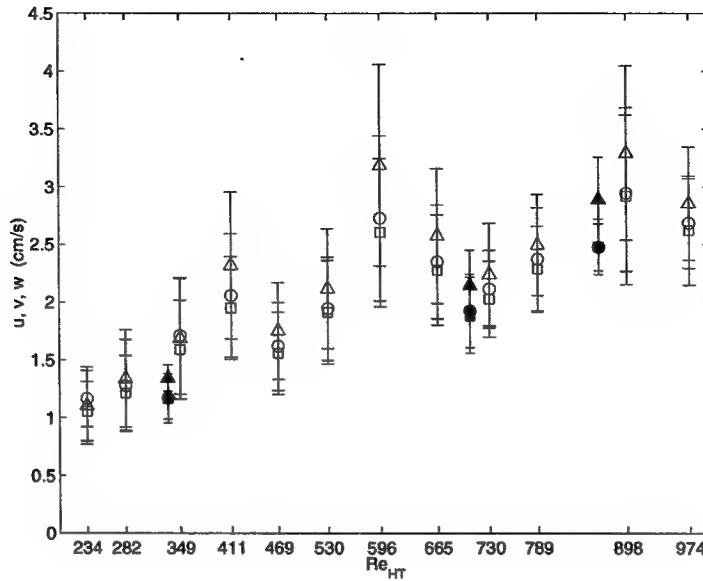


Figure 6-16: RMS turbulent velocity fluctuations as a function of grid Reynolds number. Each point corresponds to the spatial average over the nine measurement locations, and the vertical bars indicate plus/minus one standard deviation. Squares correspond to RMS u velocities, circles correspond to RMS v velocities, and triangles correspond to RMS w velocities. The solid markers indicate grid B results; all others, grid A.

larger than those predicted, and a relation using $C_1 \approx 0.35$ yields better agreement.

Although implicit in some of the previous figures, the horizontal spatial variability of the turbulence is made more clear in figure 6-19. This figure shows the relative standard deviation of the RMS turbulent velocities over the Reynolds numbers studied. As noted earlier, the grid B flows appear to exhibit the most homogeneous behavior. The largest degree of *absolute* variability was found for the highest oscillation frequencies, $Re = 596$ and $Re = 898$, where $f \geq 2.2$ Hz. Overall, the relative spatial variability of the RMS turbulent velocity components is approximately 20%. Over all Reynolds numbers, this corresponds to an actual variability of roughly 0.5 cm/s. As shown earlier, this variability derives mostly from the disparate nature of the corner region flows.

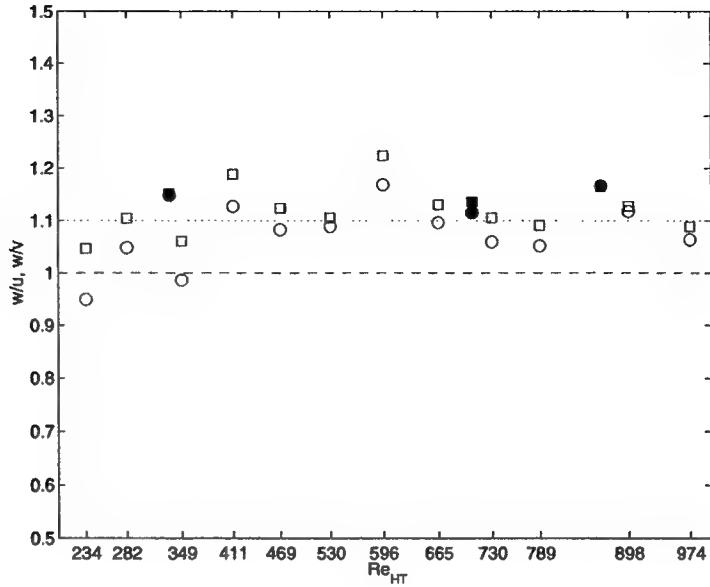


Figure 6-17: Isotropy of the grid-generated turbulence as a function of grid Reynolds number. The points reflect the ratios, w/u (squares) and w/v (circles) of the spatially averaged values of the fluctuating velocity components. The dashed line indicates the result for perfect isotropy. The dotted line represents the estimated ratio for the actual data ($C_2 \cong 1.1$). The solid markers indicate grid B results; all others, grid A.

6.3.7 Discussion

These ADV velocity measurements have provided a number of useful results. First, they have indicated that the grid tank flows are relatively repeatable. Worst case estimates would predict run-to-run variability of 15–20%. As mentioned, such variability is not too critical for this thesis work, as the experimental measurements of chapter 7 are made such that the results can be considered on an individual run basis. The ADV results have also shown that the bulk flow can be treated as stationary. Additionally, these measurements have quantified some of the actual turbulence levels in the grid tank. In doing so, these results have provided further verification of the empirical relations of Hopfinger and Toly. This lends confidence to the use of these relationships for estimating the turbulence elsewhere in the tank (e.g., other depths), and for additional grid forcing conditions. The ADV results have also shed light on the nature of the mean flows that are generated in the grid tank. These flows are not insignificant, although they are typically relatively small fractions of the

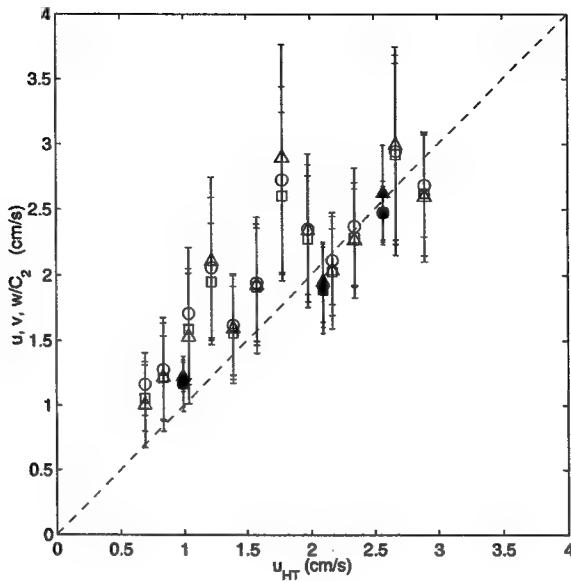


Figure 6-18: Comparison between empirically predicted RMS turbulent velocities and those measured. Each point corresponds to the spatial average over the nine measurement locations, and the vertical bars indicate plus/minus one standard deviation. Squares correspond to RMS u velocities, circles correspond to RMS v velocities, and triangles correspond to scaled RMS w velocities. The solid markers indicate grid B results; all others, grid A. The dashed line is the 1:1 relationship.

turbulence levels. Further discussion of the mean flows is left for the summary in section 5.4. Finally, the velocity measurements with the ADV have revealed that the flow, both mean and turbulent, is not ideally homogeneous in horizontal space as would have been hoped. While the degree of inhomogeneity is not large, this fact could not be overlooked. The next section explores the mean flow structure, as well as the turbulence, in greater detail.

6.4 PIV measurements of bulk flow structure

To further explore the nature of the bulk flow structure and the spatial distribution of the turbulence for the grid-stirred tank, a series of whole-field PIV measurements were made for two grid forcing conditions. Measurements were made in three separate vertical planes above the grid to provide an overall assessment of the flow.

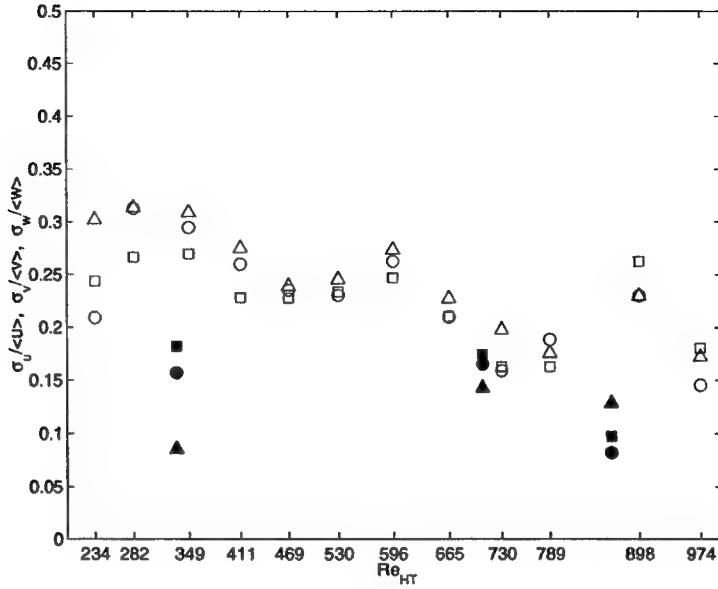


Figure 6-19: Horizontal spatial variability of RMS turbulent velocities over all Reynolds number studied. The variability is measured as the standard deviation of the nine spatial measurements for each RMS turbulent velocity component (σ_{u_i}), scaled by the mean value of the respective RMS component over the same nine locations ($\langle u_i \rangle$). Squares correspond to u components, circles correspond to v components velocities, and triangles correspond to w components. The solid markers indicate grid B results; all others, grid A.

6.4.1 Experimental arrangement and procedure

The PIV system discussed in section 4.4.3 was utilized in a standard light sheet mode in order to map two-dimensional planes of the flow at a number of locations in the grid tank. The camera was mounted to a two-dimensional, horizontally translating system that allowed the camera to image a range of planes in the tank, and repeatably return to these locations. For these measurements, the tank was cleaned, filled, and seeded with the fluorescent particles. No optical filtering of the camera took place. The camera timing and image acquisition were set up such that the camera operated in double exposure mode, acquiring a pair of video frames every 3 seconds for approximately 8 minutes (156 frame pairs). The time between frame exposures, Δt , was 10 ms throughout these experiments. For both of the grid forcings, the grid was run for a half hour before measurements were taken. An 8 minute sequence was collected at each of the nine selected image areas. Three

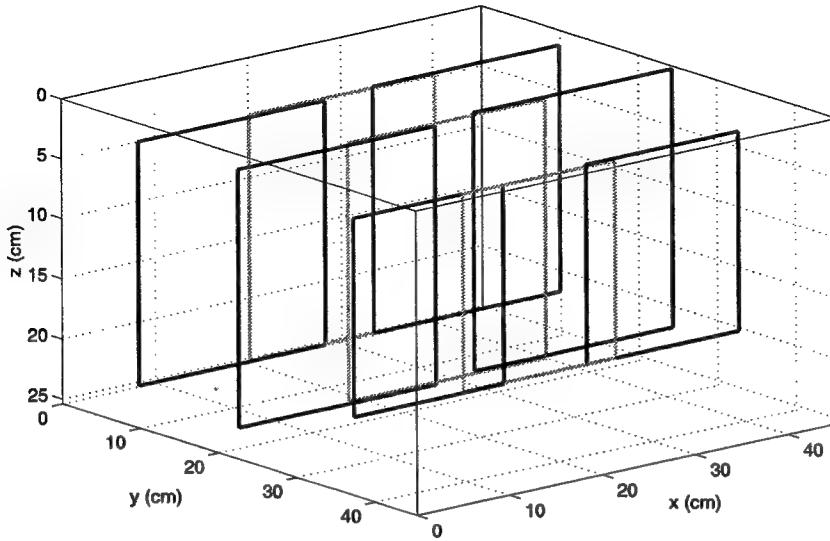


Figure 6-20: Location of the imaged areas for the bulk flow PIV mapping measurements. The three-dimensional domain corresponds to the physical domain of grid tank with the lower z -plane located at the mean grid position ($z = 25.4$ cm). The nine measurement planes are shown. The middle area in each plane is outlined in gray to improve clarity. The planes are located at $y = 10, 23$, and 35 cm.

slightly overlapping areas were taken in a single xz -plane, and three separate xz -planes were investigated. The image areas are shown in figure 6-20. These planes correspond to the same planes that the earlier ADV measurements were located in: $y = 10, 23$, and 35 cm. The collection of imaged areas essentially span the complete x -direction of the grid tank, approach the free surface to within 0.7–3.1 cm, and measure down to near the top of the grid stroke. The light sheet thickness was made to be approximately 3 mm. This relatively large thickness was chosen in order to retain particle images between frame pairs in this highly three-dimensional flow.

The two conditions studied correspond to $Re = 596$ and $Re = 730$, both with grid A. The two were selected as representative conditions in terms of mean flow behavior. The earlier ADV measurements indicated that strong, localized jets existed for the $Re = 596$ forcing; the $Re = 730$ represented a potentially different regime where such jets were not as predominant and the vertical mean flow appeared more spatially homogeneous. Results from these two conditions were expected to reveal differences between the case of strong jet

flows and the case of weaker jet flows.

The PIV image data were processed identically for all nine locations, and both grid conditions. The hybrid correlation scheme was used, employing subimages 32 px on a side. Since the image physical calibration varied from plane to plane, the physical dimensions of the PIV interrogation window also varied; for the $y = 10$ cm plane, the subimage corresponded to 6.4 mm square, for the $y = 23$ cm plane, the subimage corresponded to 6.8 mm square, and for the $y = 35$ cm plane, the subimage corresponded to 5.2 mm square. Vectors were computed every 16 px in both directions. The physical spacing of vectors can be deduced from the above specified dimensions. Finally, a surface tension measurement at the conclusion of the experiment indicated a final surface pressure of -0.10 ± 0.28 mN/m.

6.4.2 Bulk PIV results for $Re = 596$

The goal of these experiments and their analysis was to adequately map the flow and investigate the basic properties of the turbulent field. In that vein, spatial mappings of the mean flow and turbulent energy, RMS turbulent velocities, and integral lengthscales were investigated.

Mapping the mean flow and turbulent kinetic energy

Complete two-dimensional pictures of the mean flow and turbulent kinetic energy were produced for each of the three xz -planes. The turbulent kinetic energy (per unit mass) was defined as

$$q(x, z; y) = \frac{1}{2}(2u^2 + w^2), \quad (6.9)$$

where the assumption of horizontal isotropy, $u = v$, has been used. For a single xz -plane, the 156 velocity field measurements at each imaged area were ensemble-averaged to produce a mean velocity field: $U(x, z; y), W(x, z; y)$. This mean velocity field was subtracted from each individual velocity field realization for that area to yield the turbulent fluctuation velocity fields: $u(x, z; y), w(x, z; y)$. The turbulent fluctuation velocity fields were ensemble-averaged to produce a mean turbulent kinetic energy field. The three mean velocity fields were assembled into a single two-dimensional data set and mathematically interpolated to

a full xz -plane that spanned the complete x -direction of the grid tank. The same procedure was performed for the turbulent kinetic energy. Figure 6-21 shows the results from this mosaicing process for $Re = 596$. The regions where $x > 20$ cm and no measurements exist are due to the absence of light sheet illumination. Inspecting first the mean flow behavior, the PIV results provide exceptional whole-field quantitative visualizations of the persistent bulk flow structure. The results from the three planes confirm the hypothesized mean circulations based on the spatially coarse ADV measurements. Evidence of consistent upwelling is found in the majority of the tank (above the grid) with a highly localized region of downwelling located in the center region of the tank. Significant additional downwelling occurs along all side walls. In response to this flow pattern, a number of rotational cells emerge in the mean. One such cell is obvious in the $y = 10$ cm plane, and another is present in the opposite corner of the tank in the $y = 35$ cm plane. A small mean roll cell also appears to exist near the free surface in the $y = 10$ cm plane, generated by the converging near-surface flows whose origins are the corner jets. Some flow symmetry can be found from these mappings. It appears that the mean flow may show symmetry about one of the tank diagonals (expressed by the line $y = 45.4 - x$). It also might be speculated that the two tank vertical midplanes exhibit some degree of symmetry.

Turning to the turbulent kinetic energy, the spatial variability is somewhat less dramatic. In the $y = 10$ cm plane and the $y = 35$ cm plane, the turbulent energy is roughly constant across the x -direction for a given depth. However, higher intensities are found in the corner regions. In the midplane ($y = 23$ cm), the turbulent energy undergoes a significant reduction in the center region where the mean downwelling occurs. This would indicate that for this grid forcing, the central region of the flow is a relatively low energy region with a persistent downward flow. To either side of this region, the net flow is upward and the turbulence levels are similar to those in other regions about the tank.

RMS turbulent velocities

The vertical decay of the RMS turbulent velocity fluctuations was investigated for the various nine measurement areas. Both horizontal (u only) and vertical fluctuations were considered. In each area, and for each velocity component, the 156 velocity fields were

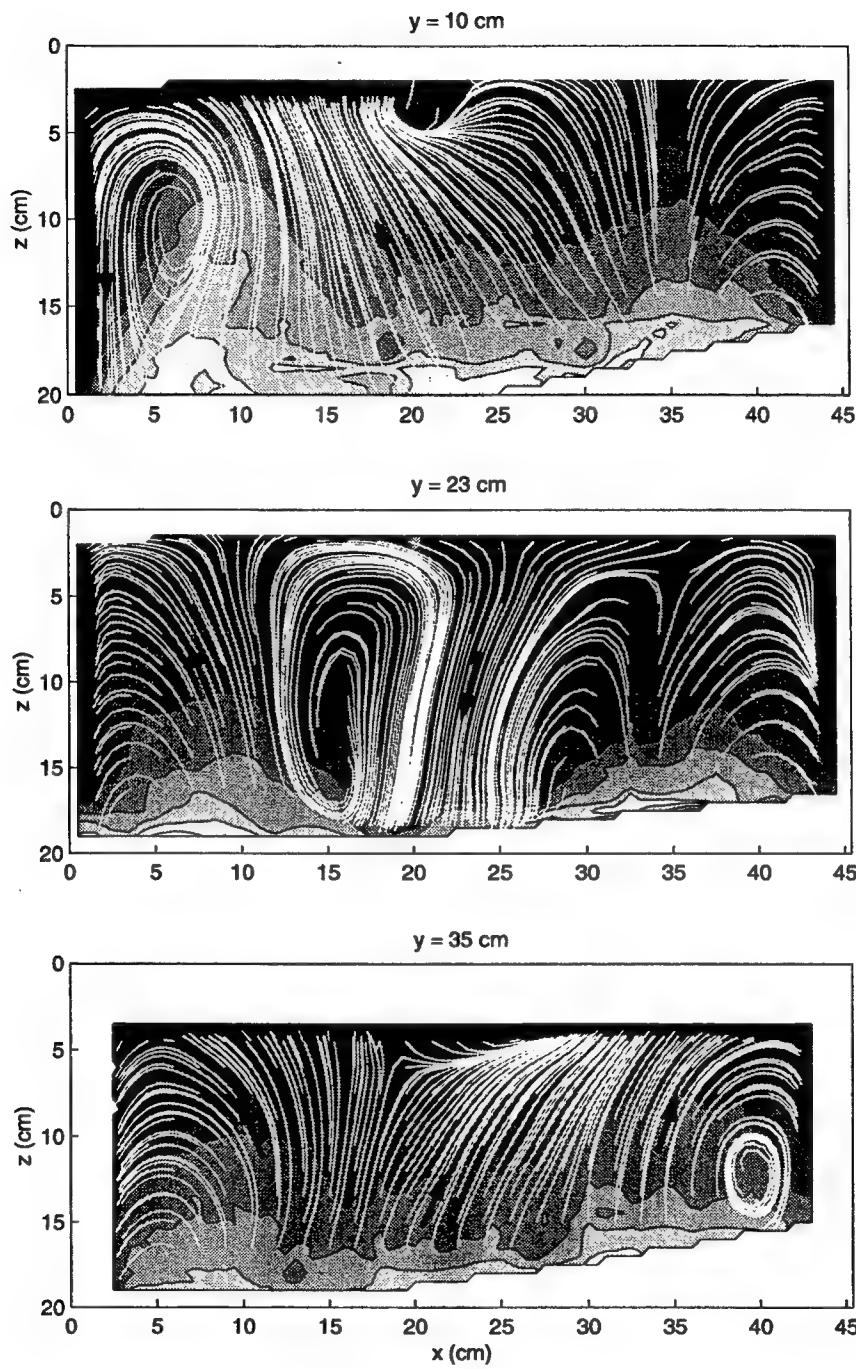


Figure 6-21: Mean flow and turbulent kinetic energy visualization in three vertical planes for $Re = 596$. The ensemble-averaged mean flow (streamlines) and turbulent kinetic energy (contours) are shown for the xz -planes at $y = 10, 23$, and 35 cm . The turbulent kinetic energy is shown as the base-10 logarithm of the actual energy for optimal visualization. The free surface is $z = 0$, the top of the grid stroke is at $z = 21.6 \text{ cm}$, and the mean grid position is $z = 25.4 \text{ cm}$. The contour levels in each panel, and those of figure 6-27 are identical.

used to generate an ensemble-averaged turbulent velocity fluctuation field. This field was averaged in the x -direction (over roughly 20 cm) to produce a single vertical profile of each velocity component at each measurement area. These results are presented in figure 6-22. In addition to the measured profiles, also shown for reference is the empirical prediction result of Hopfinger and Toly for this set of grid forcing parameters. A number of things are evident from this series of results. First, as already seen, there exists a certain degree of spatial inhomogeneity in the RMS turbulent velocities. This inhomogeneity aside, the decay law of $u \sim z^{-1}$ appears to be followed reasonably well in the different regions throughout the tank. The most apparent deviation from this relationship occurs for large z where the measurements are very near the oscillating grid, and the empirical prediction is likely not valid. Additionally, it is found that the vertical fluctuations are almost always greater than the horizontal, by an amount in range of the predicted 10–30%. Lastly, the influence of the free surface on the turbulent field is observed. It appears that as the surface is approached, the horizontal fluctuations no longer show a decay, but rather experience a slight increase; the vertical fluctuations, on the other hand, appear to show the beginning of a decrease near the surface.

To gain an overall assessment of the vertical decay in the tank as a whole, the measurements from figure 6-22 were collected to form a single profile of each velocity component for the entire tank. This was accomplished by bin-averaging the complete set of measurements into 4 mm vertical bins. The result is shown in figure 6-23. Also included in this figure are the empirical predictions u_{HT} (equation 6.1) and w_{HT} (equation 6.2), as well as numerical approximations based on the source theory of Hunt and Graham [61]. These approximations were taken from the work of Brumley and Jirka [13], who found that the source theory of Hunt and Graham is appropriate to describe the turbulent flow field in the surface-influenced layer. The source theory treats the turbulent field as a superposition of homogeneous turbulence and an irrotational velocity field driven by a source distribution at the interface that causes the vertical velocity fluctuations to vanish there. Velocity profiles based solely on the Hunt–Graham theory, and combined Hunt–Graham/Hopfinger–Toly profiles are shown.

The agreement between the measured horizontal RMS turbulent velocity profile and

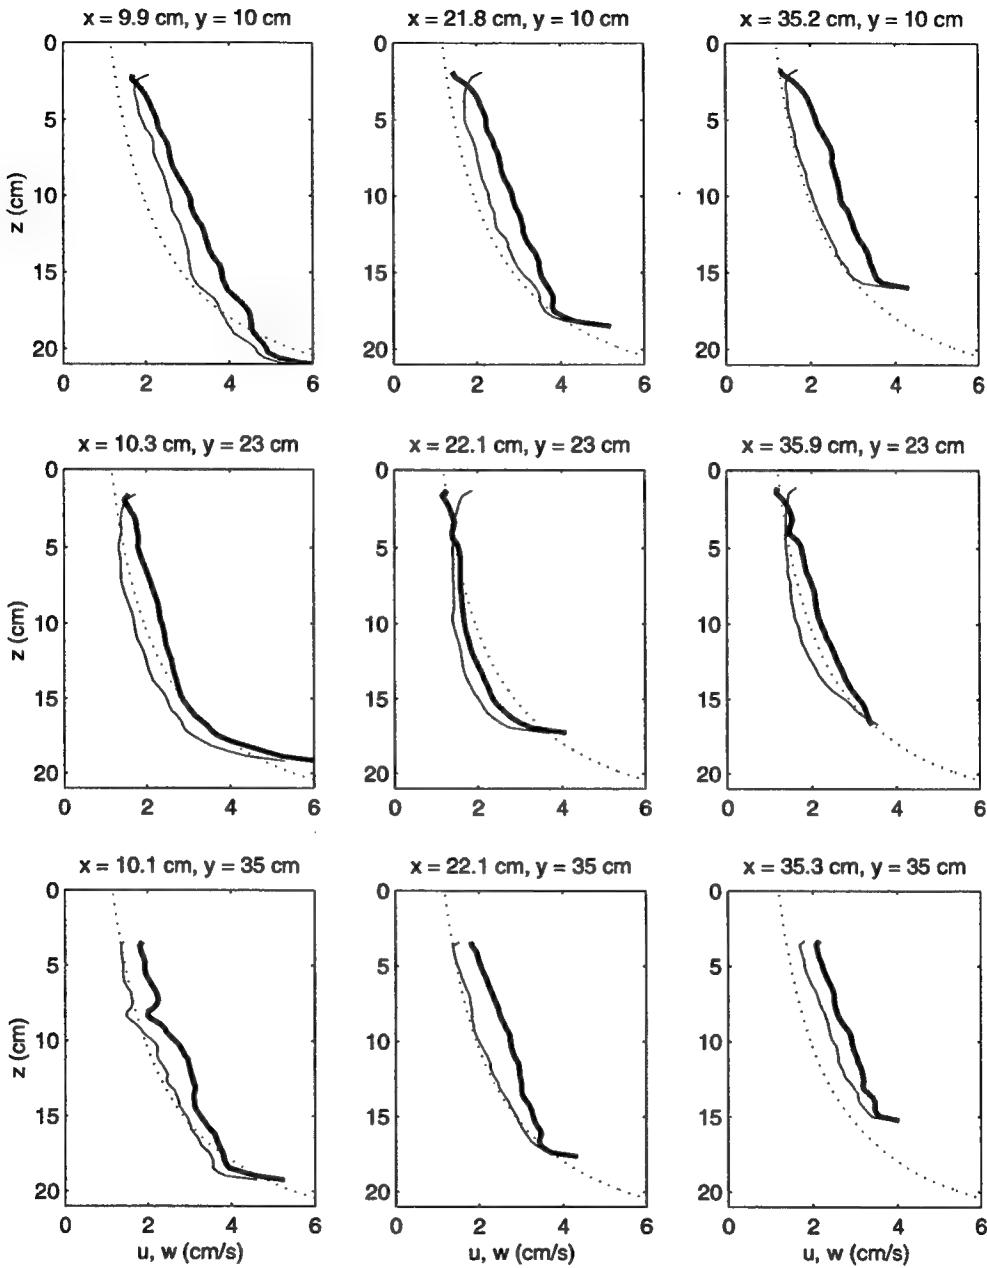


Figure 6-22: Vertical decay of the RMS turbulent velocities for $Re = 596$. Shown are the results at each measurement area. Each velocity profile is the result of the ensemble-averaged RMS turbulent velocity field that was collapsed in the x -direction to yield a single vertical profile. Shown are the horizontal turbulent fluctuations u (thin line), the vertical fluctuations w (heavy line), and the empirical prediction u_{HT} (dotted line) using $C_1 = 0.25$ for reference. Above each panel is the “mean” location of the measurement: x indicates the mean x location of the image area, and y indicates the location of the vertical xz -plane.

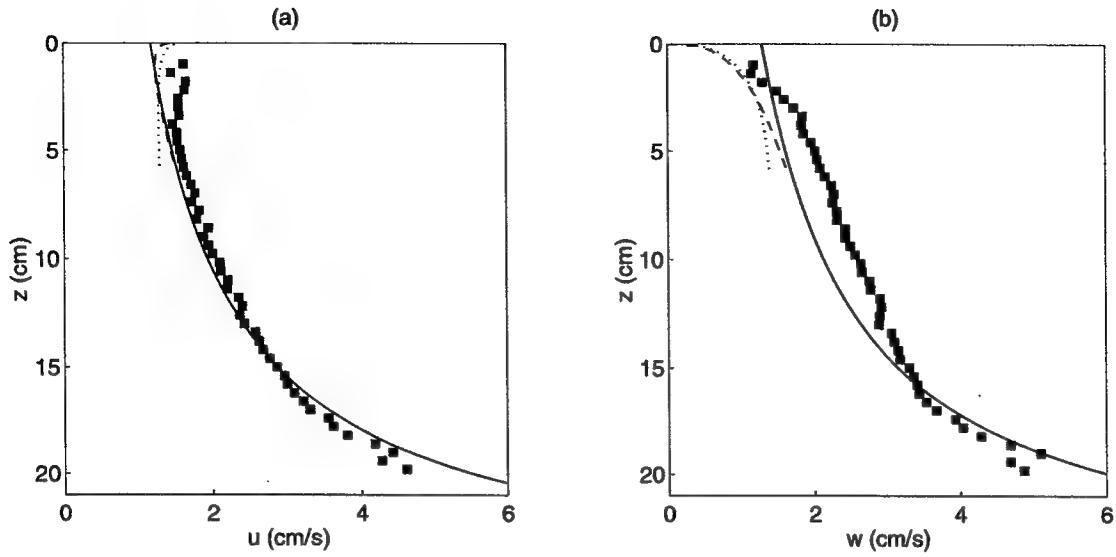


Figure 6-23: Overall mean profiles of the RMS (a) horizontal and (b) vertical turbulent fluctuations for $Re = 596$. Each profile is the result of bin-averaging (in z) the particular set of profiles in figure 6-22. Shown are the velocity fluctuation profile data (squares), the empirical predictions u_{HT} and w_{HT} (solid lines), the Hunt–Graham profiles (dotted lines), and the combined profiles (dashed lines).

that of the predicted profile of Hopfinger and Toly is marked. The effect of the free surface, based on the Hunt–Graham profile, is to marginally enhance the RMS horizontal fluctuations at the interface. It appears that the measured profile begins to show this behavior near the surface, at a depth slightly greater than that predicted from the source theory. The agreement between the Hopfinger–Toly prediction and the measured vertical fluctuations is good, but not as impressive as with the horizontal velocities. The prediction generally underestimates the magnitude of the fluctuations. The free-surface influence on the vertical velocity is dramatic. The measured vertical fluctuations show an abrupt decrease in magnitude at a depth of approximately 4 cm, closely mimicking the trend predicted by the Hunt–Graham profile. This effect is to be expected, owing to the kinematic boundary condition at the free surface that forces $w \rightarrow 0$ on $z = 0$. The depth of 4 cm is found to correspond to roughly two integral lengthscales (discussed below).

Turbulent integral lengthscales

An analysis similar to that just discussed for the velocities was also conducted to estimate the turbulent integral lengthscales. The lengthscales were calculated from the spatial auto-correlation of the fluctuating velocities along fixed values of z . The velocity two-point auto-correlation is defined as

$$R_{ij}(r) = \overline{u_i(\vec{x})u_i(\vec{x} + r\hat{e}_j)}, \quad i = 1, 2, 3; \quad j = 1, 2, 3, \quad (6.10)$$

where \hat{e}_j is the unit vector in the j -direction, and r is the distance between measurement points. Summation notation is not implied for index i . The integral scales are calculated from integration of the normalized auto-correlations:

$$L_{ij}(z; y) = \int_0^\infty \frac{\overline{u_i(\vec{x})u_i(\vec{x} + r\hat{e}_j)}}{\overline{u_i^2}} dr, \quad (6.11)$$

For these PIV measurements, the (horizontal) longitudinal integral scale is given by

$$L_{11}(z; y) = \int_0^\infty \frac{\overline{u(x)u(x + r\hat{e}_x)}}{\overline{u^2}} dr, \quad (6.12)$$

representing the correlation between horizontal velocities along the horizontal x -direction.

The (horizontal) transverse integral scale is given by

$$L_{31}(z; y) = \int_0^\infty \frac{\overline{w(x)w(x + r\hat{e}_x)}}{\overline{w^2}} dr, \quad (6.13)$$

representing the correlation between vertical velocities along the horizontal x -direction. Computed in this way, the two lengthscales represent the characteristic eddy scales at a given depth within each measurement area. Correlations were not computed vertically along fixed values of x since the lengthscales are expected to vary with z . The results for the nine measurement areas are shown in figure 6-24 for $Re = 596$.

The results for the estimated lengthscales appear to be scattered about the predicted Hopfinger-Toly relationship (equation 6.3), shown for reference. A discernible pattern to the variations is not necessarily obvious. In general, the longitudinal lengthscale increases with

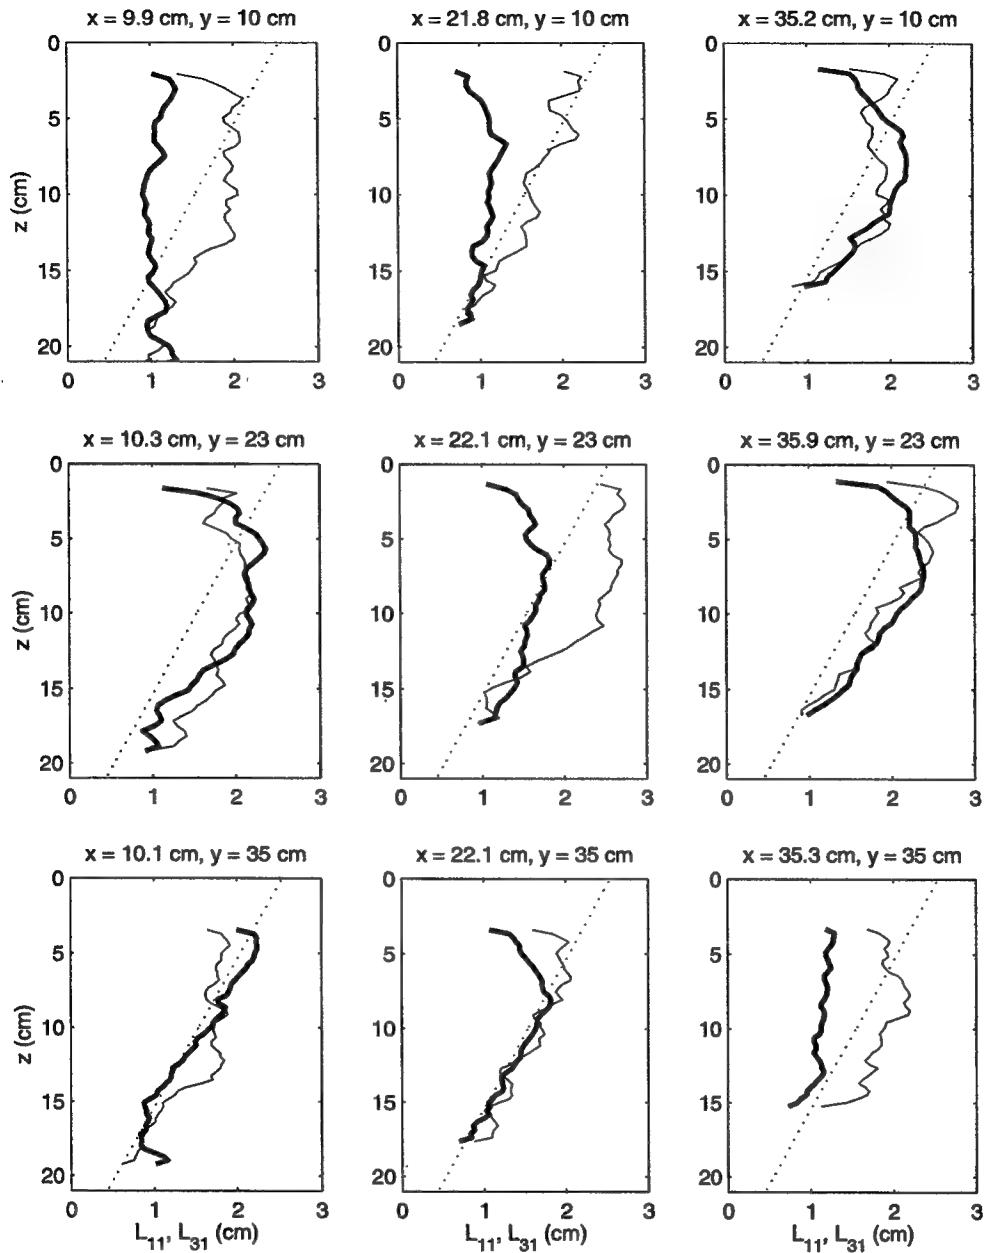


Figure 6-24: Vertical dependence of the integral lengthscales for $Re = 596$. Shown are the results at each measurement area. Each profile is the result of the ensemble average of the profiles produced from the individual velocity fields. Shown is the longitudinal lengthscale L_{11} (thin line), the transverse lengthscale L_{31} (heavy line), and the empirical prediction L_{HT} (dotted line) using $C_3 = 0.1$ for reference. The coordinates above each panel are as in figure 6-22.

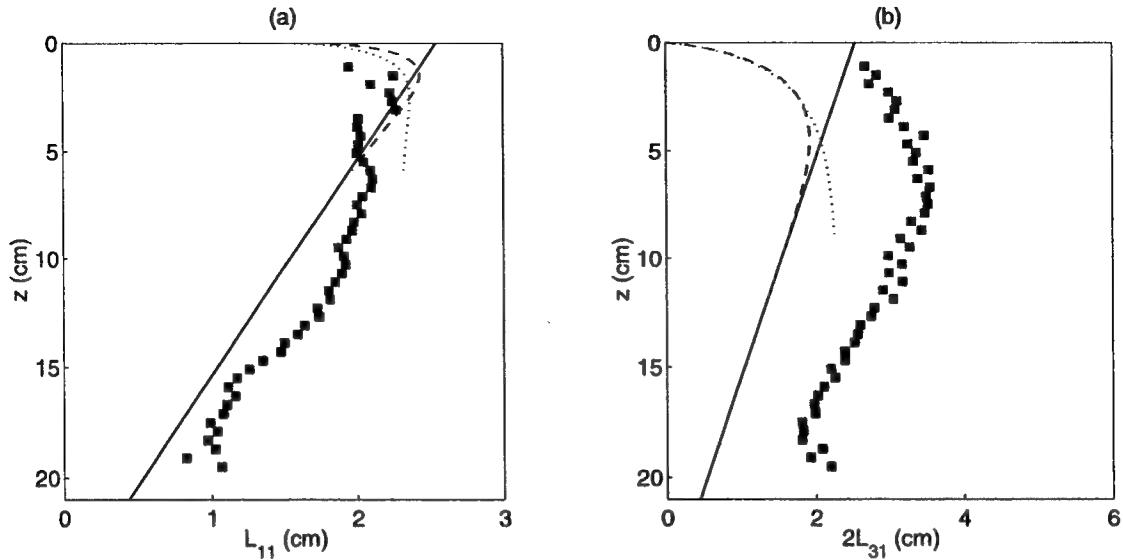


Figure 6-25: Overall mean profiles of the integral lengthscales for $Re = 596$. Each profile is the result of bin-averaging (in z) the particular set of profiles in figure 6-24. (a) Longitudinal lengthscale L_{11} , (b) doubled transverse lengthscale $2L_{31}$. Also shown are the empirical prediction L_{HT} (solid lines), the Hunt-Graham profiles (dotted lines), and the combined profiles (dashed lines).

distance from the grid, and the transverse scale increases for some distance and then falls off as the surface is approached. To facilitate discussion, similar to the RMS velocities analysis, the nine pairs of profiles in figure 6-24 were collected and bin-averaged into 4 mm vertical bins. The result is given in figure 6-25. Also included in this figure are the predicted profiles based on the Hopfinger-Toly relation, the Hunt-Graham source theory, and the combined Hopfinger-Toly/Hunt-Graham result (again, Brumley and Jirka [13]). In figure 6-25b, the *doubled* transverse lengthscale is plotted to evaluate to the relationship expected for isotropic turbulence, where the two lengthscales studied are related as $L_{11} = 2L_{31}$ (Tennekes and Lumley [129]). This states that the doubled transverse scale should be equal to the lengthscale predicted by the Hopfinger-Toly expression.

In figure 6-25a, the Hopfinger-Toly relation for the longitudinal lengthscale fits reasonably well to the overall mean L_{11} profile away from the surface ($z > 6$ cm). However, a noticeable underestimation is evident. Close to the surface, the longitudinal lengthscale exhibits a slight increase near $z = 3$ cm and appears to show the beginning of a reduction

very near the interface. This reduction agrees nicely with the source theory prediction of Hunt and Graham, and similar behavior was reported by Brumley and Jirka. The transverse lengthscale follows the same trend as the longitudinal lengthscale for depths greater than 6 cm, and is consistently greater than the predicted values by about a factor of two. For $z < 6$ cm the transverse lengthscale begins to decrease as the surface is approached, following the trend predicted by the Hunt–Graham theory. This is a direct consequence of the free-surface boundary condition previously noted. Physically, these scales provide a sense of the shapes and sizes of the eddies comprising the grid-generated turbulence. The effect of the free surface is to introduce a “source” or “blockage” layer, which is a geometric consequence of eddies approaching the surface. This layer gives rise to a “pancake” eddy structure (Handler *et al.* [47]) near the interface where the vertical extent of the eddies is reduced and the lateral extent changes little. Similar results for the Taylor microscales were found in the direct numerical simulation work of Shen *et al.* [120].

The isotropic turbulence relation, $L_{11} = 2L_{31}$, is not found to hold true for the present measurements of these lengthscales. Instead, these data indicate $L_{11} \approx 1.2L_{31}$. If the longitudinal lengthscale is assumed a fair estimate, this would imply that the transverse scale exceeds its predicted isotropic value by nearly a factor of two. It is not clear if this result is a consequence of a lack of flow isotropy. In similar grid tank experiments, Kit *et al.* [71] were only able to compute L_{31} , and the results for this lengthscale appeared to follow the relationship for L_{HT} , which is for the *longitudinal* lengthscale, *i.e.*, $L_{11} \approx L_{31}$. This also would imply the transverse scale is larger than the isotropic prediction. Similarly, Brumley and Jirka [13] find that $2L_{31}$, which should be equal to the longitudinal scale result, is underpredicted by the Hopfinger–Toly longitudinal estimate in the bulk. However, they do find that $L_{11} \approx 2L_{13}$ away from the free surface. It is possible that, for the present measurements, the PIV data field is not sufficiently expansive and/or resolved to yield accurate lengthscale estimates.

6.4.3 Bulk PIV results for $Re = 730$

The same collection of measurements and results for the grid forcing condition with $Re = 730$ are presented here. Since the analysis was identical to that for $Re = 596$, details are omitted

and the results simply are presented and discussed.

Mapping the mean flow and turbulent kinetic energy

Figure 6-26 shows the mapping results for the mean flow and turbulent kinetic energy for this Reynolds number. Consistent with the ADV results, the mean flow for this grid forcing does not show strong localized jet structures as in the case for $Re = 596$. Instead, the mean flow is upward throughout the tank, with the downward return flow existing exclusively at the tank side walls. Unlike the $Re = 596$ condition, mean rotational cells do not appear to form in the bulk for this flow, with the exception of a weak circulation near the free surface in the $y = 10$ cm plane. Such a cell also was detected for the $Re = 596$ flow. In addition, it appears the same diagonal symmetry that was observed for the $Re = 596$ case is also found for this grid forcing. This is evident from the mean streamlines for the $y = 10$ cm and the $y = 35$ cm planes. Considering the turbulent kinetic energy, the distribution of the energy is much more spatially homogeneous as compared to the $Re = 596$ case. The midplane shows slightly less turbulent energy than the other two planes, and the center region of the tank is again the area of weakest turbulent intensity.

RMS turbulent velocities

The vertical decay behavior of the RMS turbulent velocities for this Reynolds number is shown in figure 6-27. The profiles are qualitatively similar to those found for the $Re = 596$ case, but the Hopfinger–Toly (Hunt–Graham) curves consistently overpredict the observed results for this forcing condition. Such deviation is not necessarily surprising, as similar results were shown by Brumley and Jirka [13]. The shape of the profiles is in very good agreement with the predicted z^{-1} decay law. Some diagonal symmetry is also indicated by these turbulent velocity profile results. As done for $Re = 596$, these nine pairs of profiles were bin-averaged to produce overall mean profiles of RMS turbulent velocities. These profiles are given in figure 6-28 and provide a better basis for comparing the measurements with the predicted curves. The surface-influenced layer appears to exhibit behavior similar to the $Re = 596$ result. Again, very near the interface, the horizontal turbulent fluctuations show a slight increase, in agreement with the source theory prediction. Like the $Re = 596$ result,

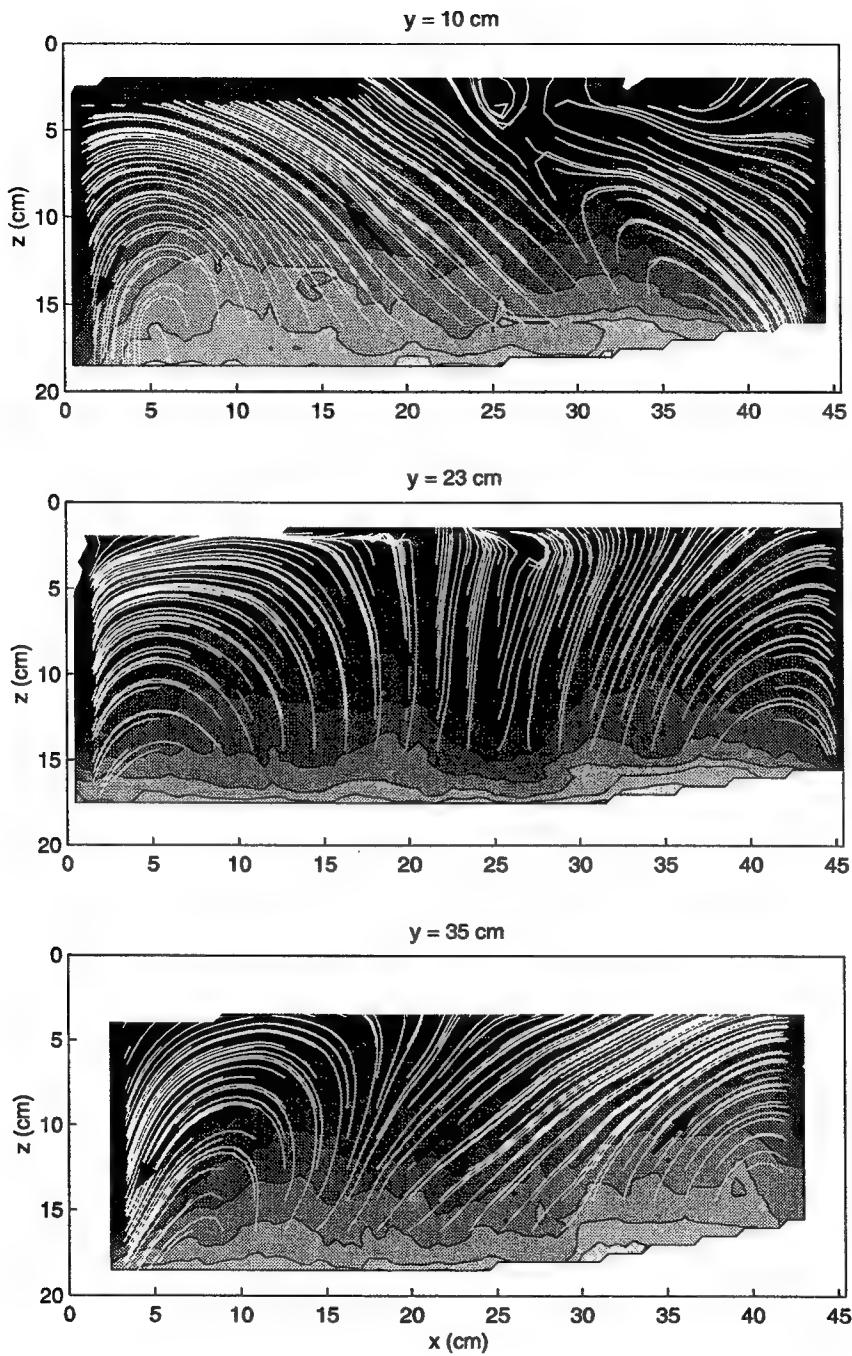


Figure 6-26: Mean flow and turbulent kinetic energy visualization in three vertical planes for $Re = 730$. The ensemble-averaged mean flow (streamlines) and the turbulent kinetic energy (contours) are shown for the xz -planes at $y = 10, 23$, and 35 cm. The turbulent kinetic energy is shown as base-10 logarithm of the actual energy for optimal visualization. The free surface is $z = 0$, the top of the grid stroke is at $z = 19.7$ cm, and the mean grid position is $z = 25.4$ cm.

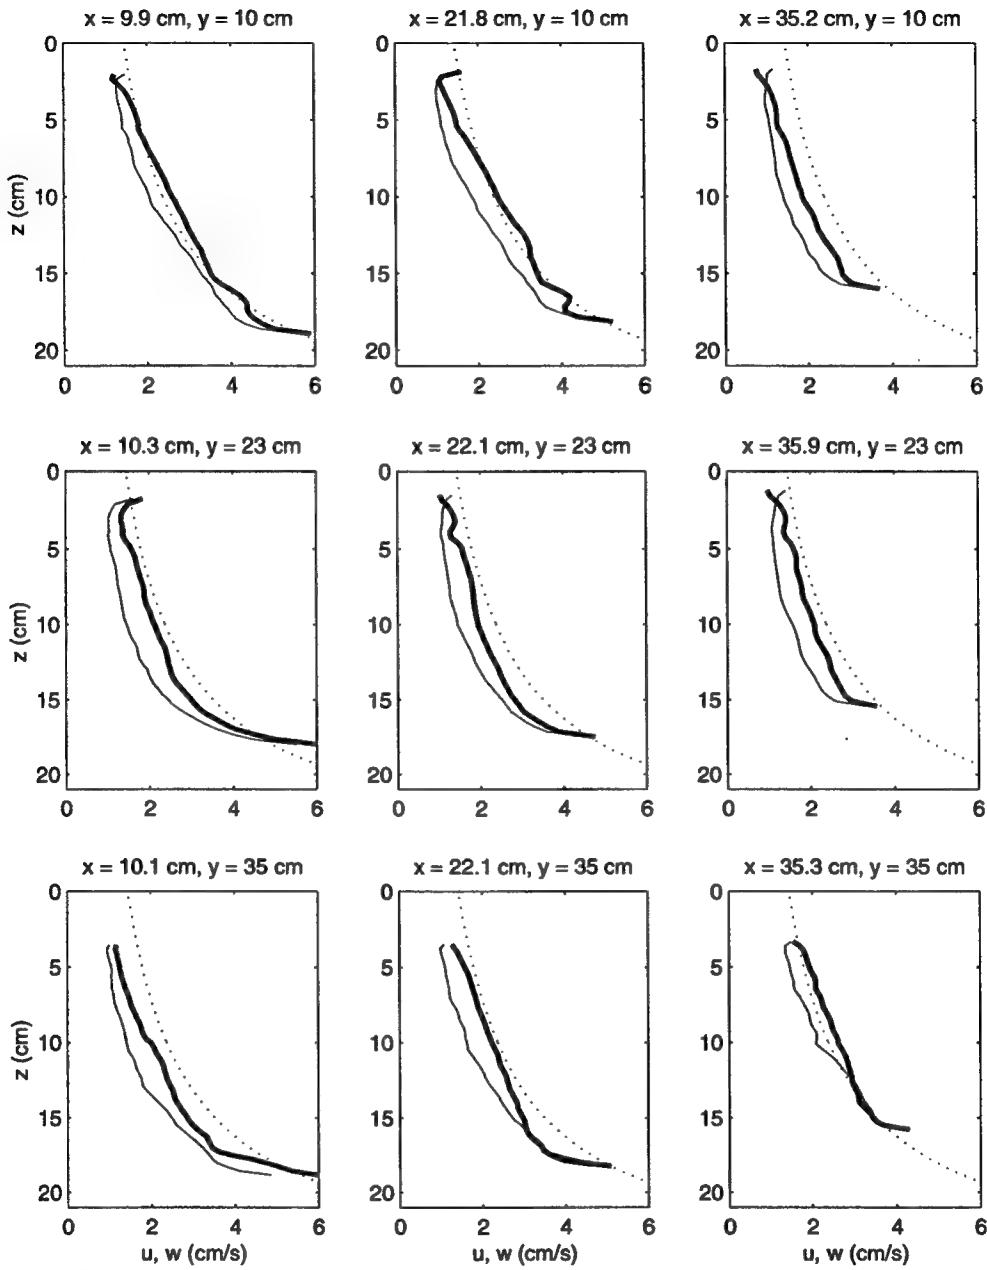


Figure 6-27: Vertical decay of the RMS turbulent velocities for $Re = 730$. Shown are the results at each measurement area. Each velocity profile is the result of the ensemble-averaged RMS turbulent velocity field that was collapsed in the x -direction to yield a single vertical profile. Shown are the horizontal turbulent fluctuations u (thin line), the vertical fluctuations w (heavy line), and the empirical prediction u_{HT} (dotted line) using $C_1 = 0.25$ for reference. The coordinates above each panel are as in figure 6-22.

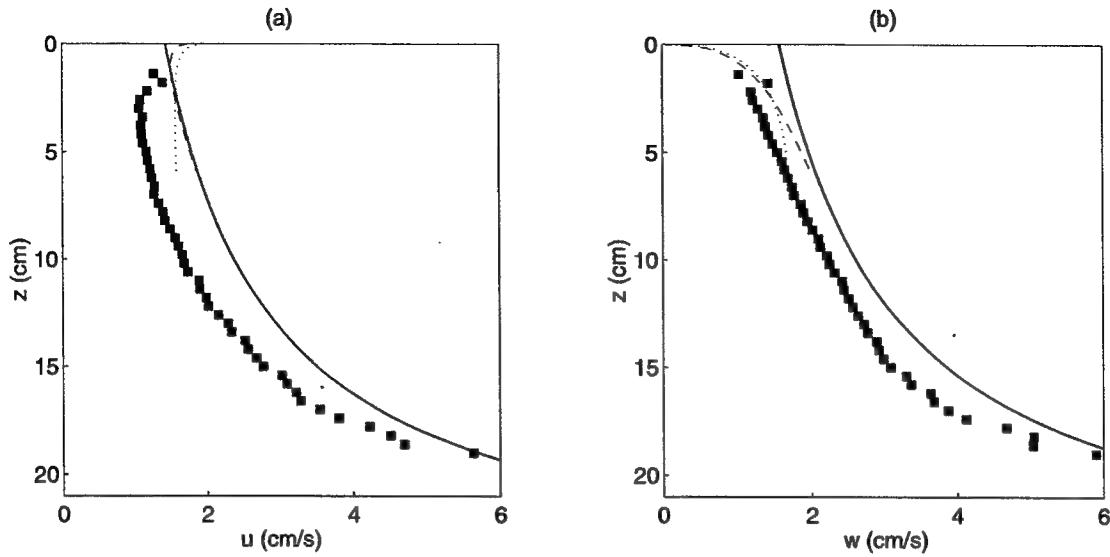


Figure 6-28: Overall mean profiles of the RMS (a) horizontal and (b) vertical turbulent fluctuations for $Re = 730$. Each profile is the result of bin-averaging (in z) the particular set of profiles in figure 6-27. Shown are the velocity fluctuation profile data (squares), the empirical predictions u_{HT} and w_{HT} (solid lines), the Hunt–Graham profiles (dotted lines), and the combined profiles (dashed lines).

this enhancement occurs somewhat deeper than what the Hunt–Graham theory predicts. A sharp reduction in the fluctuating vertical velocity component is not explicitly detected, and it is likely that this modulation occurs for $z < 1.5$ cm.

Turbulent integral lengthscales

The turbulent integral lengthscales for $Re = 730$ are presented for the nine measurement areas in figure 6-29. As with the $Re = 596$ data, there is a large degree of scatter in these lengthscale profiles. The overall bin-averaged profiles are shown in figure 6-30. The averaged profiles of the lengthscales appear very similar to those found for the $Re = 596$ forcing, confirming the L_{HT} empirical prediction that does not depend on the grid geometry or oscillation parameters. Figure 6-30a shows that, with some scatter, the longitudinal lengthscales undergo a reduction near the free surface. In figure 6-30b, the shape of the transverse lengthscales profile is in excellent agreement with the predicted curves. However, it is again roughly twice its predicted isotropic value. The reason for the deviation of the

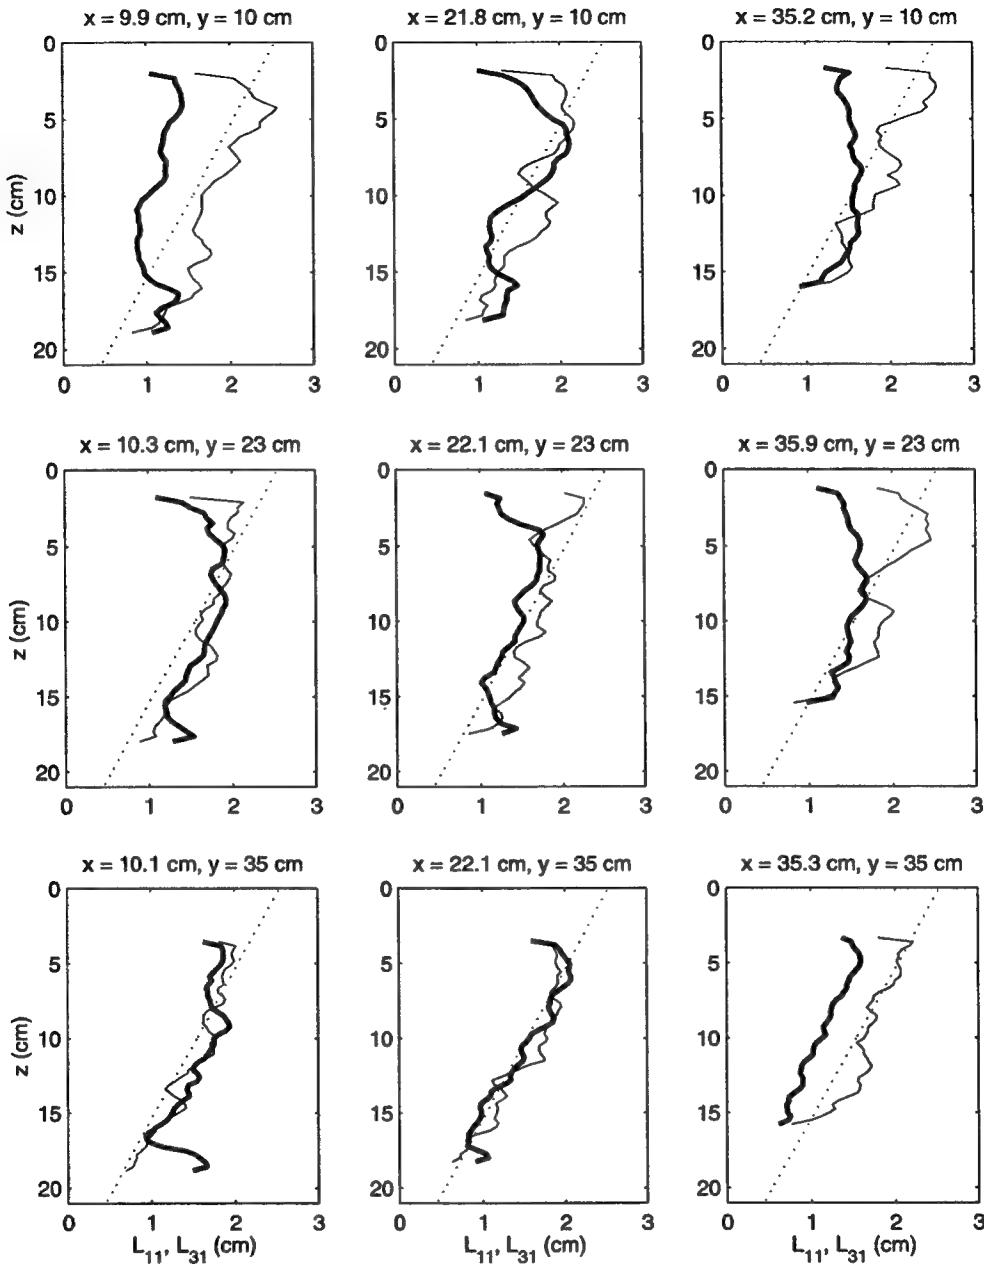


Figure 6-29: Vertical dependence of the integral lengthscales for $Re = 730$. Shown are the results at each measurement area. Each profile is the result of the ensemble average of the profiles produced from the individual velocity fields. Shown is the longitudinal lengthscale L_{11} (thin line), the transverse lengthscale L_{31} (heavy line), and the empirical prediction L_{HT} (dotted line) using $C_3 = 0.1$ for reference. The coordinates above each panel are as in figure 6-22.

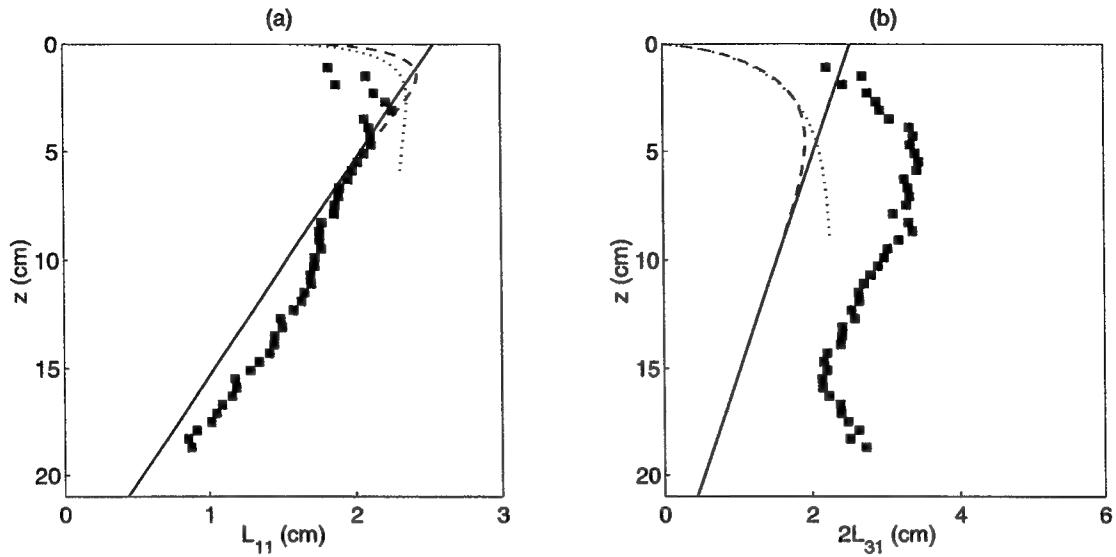


Figure 6-30: Overall mean profiles of the integral lengthscales for $Re = 730$. Each profile is the result of bin-averaging (in z) the particular set of profiles in figure 6-29. (a) Longitudinal lengthscale L_{11} , (b) doubled transverse lengthscale $2L_{31}$. Also shown are the empirical prediction L_{HT} (solid lines), the Hunt–Graham profiles (dotted lines), and the combined profiles (dashed lines).

transverse scale from both the empirical prediction and the longitudinal scale near the grid is unclear.

6.4.4 Discussion

Although only two Reynolds numbers were studied intensively using PIV, the results have provided a clear picture of the overall nature of the flow within the grid turbulence tank. The flow mapping mosaics beautifully show the complete structure of the mean circulations in the grid tank. These maps also show the spatial distribution of the turbulent energy across the tank. Since the two Reynolds numbers selected were felt to span a possible range of flow patterns, the results presented here may serve as qualitative bounds for the complete group of forcing conditions. It appears that as with the ADV measurements, the empirical relations of Hopfinger and Toly are effective predictors of the turbulent velocities throughout the majority of the grid tank. The source theory of Hunt and Graham appears to accurately predict the influence of the free surface on both RMS turbulent velocity components—the

horizontal component is slightly increased at the surface and the vertical component is forced to zero. This implies a redistribution of kinetic energy from the vertical component to the horizontal components near the interface. The Hopfinger–Toly empirical lengthscale estimate matches the measured results reasonably well away from the interface, although the measured transverse scale was found to be roughly twice its expected isotropic value. The trends for the lengthscales near the free surface predicted with the Hunt–Graham theory and the trends exhibited by the measurements agree nicely. The effect of the free surface is to reduce the horizontal, longitudinal lengthscale and drive the vertical transverse scale to zero at the surface.

Originally, it had been hoped that the grid-generated turbulence would be homogeneous throughout the tank, particularly in planes parallel to the grid. This would allow surface PIV measurements to be made at an arbitrary *localized* position at the air–water interface and be considered essentially reflective of the entire free surface. However, in light of both the ADV and PIV results, this approach had to be altered. To what degree the free surface was inhomogeneous became the next question. An answer was sought that is described in the following section.

6.5 Infrared imagery of free-surface features

Working with the grid-stirred tank, it became apparent that the free-surface flow could not be considered spatially homogeneous. This was evident from the results of both the ADV and the PIV measurements, as well as from visual observation. Often, the presence of the corner jets could be detected at the free surface through the appearance of surface boils. To quantify the surface flow variability, surface PIV was considered as an option. However, PIV was deemed impractical. The measurement domain for the surface PIV setup was only about 8 cm square. This would require roughly 16 measurements to obtain a complete picture of the surface flow for a *single* grid forcing. Investigating all twelve Reynolds numbers may not have been necessary, but half that number would have necessitated almost 100 PIV sequences, which was thought prohibitively time consuming, not to mention an image storage debacle. An alternative measurement technology therefore was considered.

6.5.1 The infrared imager and experimental arrangement

For a short period of time, this author was able to secure the use of a high performance infrared (IR) imaging system. Such systems have been recently utilized by a number of groups in the laboratory and the field to study free-surface renewal processes, in a similar vein to the efforts in this thesis work (e.g., Haussecker *et al.* [50]; McKeown and Leighton [97]; Zappa [145]). The basis for the technique lies with thermal absorption and emission in water. The absorption/emission of radiation in water is wavelength dependent. Infrared radiation propagates at wavelengths on the order of 1–10 μm . This wavelength is strongly absorbed by water, only penetrating to depths on the order of 100 μm . In a complementary fashion, the radiation at this wavelength is emitted from the upper 100 μm of the water surface. Thus, IR imagers typically will only detect the radiation that is emitted from the upper 100 μm of fluid at the air–water interface. The depth to which the imager detects is commonly referred to as the optical depth.

This absorption/emission behavior can be exploited to visualize thermal structure at the air–water interface. Typically, a “cool skin” will exist at an air–water interface owing to evaporation at the surface. Fluid at the very surface, therefore, has a slightly lower temperature than the fluid beneath it. This discrepancy can be enhanced by heating the bulk fluid beyond the air–water equilibrium temperature. Therefore, when a dynamical processes acts to bring fluid from below to the surface, this is revealed as a local temperature increase in the surface patch that has been disturbed. An IR imager senses these *surface* temperature differences as thermal radiation differences. In this way, surface renewal processes and microscale breaking at the air–water interface, for example, can be explored. As such, the IR technique can be thought of as similar to a dye flow visualization technique where the tracer is simply the fluid temperature. It is therefore a completely non-invasive, powerful method that can produce highly informative pictures of the thermal structure at an air–water interface. The relevant benefit in the present case was the fact that the camera could image a very large region of the tank water surface at once. Thus, only one measurement was required, rather than 16, and a large number of grid forcing conditions could be explored easily and efficiently.

The imager used for these experiments employed an indium antimonide (InSb) detector

with a spectral range in the mid-IR of 1.0–5.4 μm . This wavelength sensitivity gave a mean optical depth of approximately 10 μm . The focal plane pixel array of the device was 320×256 and the light integration time was 16 ms. The camera output was 12-bit digital video in a 320×240 display format at 60 Hz. The video was acquired in real-time through a digital frame capture board that was integral to the IR system computer. The imager was mounted above the tank, a distance of almost 1 m from the water surface. The imager was oriented at an angle of 37° from vertical and slightly offset to avoid spurious reflections and a reflected image of the camera itself. The tank cover was removed as the camera cannot image through the polycarbonate. The air temperature over the course of the experiments was nominally 23°C and was monitored during each run. To provide for significant free-surface temperature fluctuations, the tank spring water was heated (in the containers) in a thermal testing chamber adjacent to the laboratory. The water was brought up to a temperature of approximately 45°C and then added to the grid tank. This procedure avoided any possibility of introducing contamination to the water by alternative heating processes. A commercial probe that measured temperature, pH, and dissolved oxygen was installed in the grid tank. This probe was used throughout this thesis work for measuring temperature and dissolved oxygen time series at 0.1 Hz. For all experiments, the sensor was mounted in the tank floor using one of the sensor ports. It was located in one corner of the tank at coordinates $x = 11.4$ cm, $y = 11.4$ cm. Further details of the probe can be found in section 7.5.2. In this case, the probe was used to provide an average bulk temperature over the course of each run.

The image acquisition system was configured to collect data in two modes. First, for a long time record, the system was set up to collect one image every 2 seconds for a total of 30 minutes (900 images). This was chosen in order to study persistent surface flow features. The second mode was a visualization, burst mode. For this, the system was set up to acquire continuous image data at 30 Hz for 5 seconds (150 images). This enabled short flow sequences to be inspected for the real-time development of the flow.

For these surface visualization experiments, all of the grid conditions were studied using grid B, and $Re = 234, 282, 469, 596, 730, 898$, and 974 were studied for grid A. The reason for these choices was as follows. Based on the ADV measurements, the grid B flows appeared

to exhibit certain favorable characteristics—more spatial homogeneity and less evidence of strong corner jet activity. Further exploration of the grid B flows therefore was desired. In terms of grid A, the Reynolds numbers selected for study were felt to form a representative set of the full group of conditions. The basis for this estimation was largely the mean flow results as shown in figures 6-10 and 6-11.

Before each grid condition run, the surface was cleaned using the surface aspiration technique described in section 6.2.4. Occasionally, as the tank water temperature fell, some water would be drained, and additional 40+°C water would be added to maintain a strong air–water temperature difference. Surface tension measurements were made for these experiments using the method discussed in section 2.3.1.

6.5.2 Results and analysis

The IR imager output consists of an array of 12-bit integers that indicate relative radiation intensity. The measured intensity often is sensitive to reflections, ambient lighting and other thermal sources, and camera orientation. In order to account for such potential biases in the imagery, a mean image for a complete set of runs (all grid A runs, or all grid B runs, each group of which was performed on a separate day) was determined. This group mean image was used to re-scale the individual images in each sequence to reduce the presence of apparent persistent features due to ambient biasing.

A calibration is required to relate the IR imager counts to actual temperature values. Since the goal of these experiments was flow visualization, such a calibration was not necessary. However, an approximate calibration was determined that is reasonably correct in a relative sense; *i.e.*, the *absolute* values of the temperature are not precisely correct. This provided a physical sense of the thermal variability at the water surface. The calibration was determined by correlating the maximum imager counts with the measured bulk water temperatures. This approach was based on the assumption that the warmest fluid at the surface was fluid that was very recently brought from the bulk, and thereby reflective of the bulk water temperature. The calibration result is shown in figure 6-31. The data are represented reasonably with a linear least-squares fit. This result was used to scale the raw IR images to images of temperature.

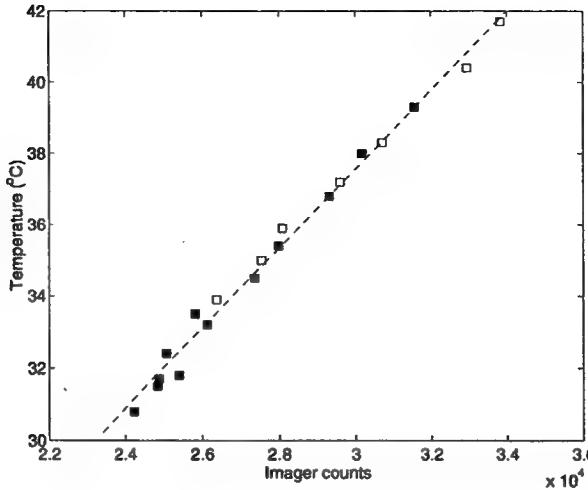


Figure 6-31: Approximate temperature calibration for the IR imager. The temperature values represent the mean bulk water temperature measured over the course of each grid forcing condition run. The imager counts reflect the maximum imager count value from the mean image of each sequence. The solid markers indicate grid B results; open markers, grid A. The dashed line indicates the linear fit result used as calibration.

Figure 6-32 shows a sample IR image from a grid A experiment at $Re = 282$. This is the raw infrared image from the camera, corrected for biasing effects, and scaled to reflect water surface temperature. Although the image is rectangular, the actual imaged area was trapezoidal in shape due to the angled, off-axis position of the IR imager. Measurements of the actual field of view allowed for the proper corrections to be made. This figure illustrates the rich behavior of the flow at the free surface. In this particular image, an area of upwelling is apparent in each of the image domain corners. This agrees nicely with the mean ADV measurements at 8 cm depth for this Reynolds number in figure 6-10.

Persistent areas of upwelling were observed frequently for all grid forcing conditions. Using the 30 minute time record, a mean thermal image was found for each image sequence. If the surface flow was ideally homogeneous and random, a mean thermal image should reveal a uniform temperature, given sufficient sampling. Deviations from such spatial uniformity would indicate persistent surface flow features. In order to present all of the results in a manner that allows for possible comparison, a dimensionless temperature was defined. For each condition, the mean thermal image was itself de-meaned and scaled by the air–water

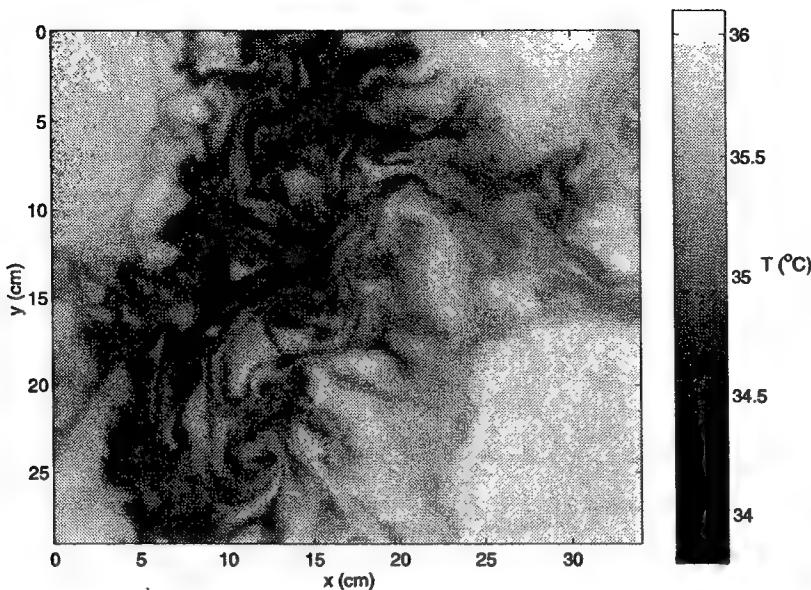


Figure 6-32: Sample IR image of the water surface for $Re = 282$ (grid A). Shown is a raw image from the camera, corrected for biasing effects, and scaled to reflect relative temperature. The axis origins are arbitrary.

temperature difference: $(T - \bar{T})/(T_w - T_a)$. Since the air-water temperature difference varied over the course of the experiments, scaling the image in this way removes the interfacial thermal gradient effect on the bulk-surface temperature difference. Therefore, each mean image illustrates the spatial variability of water surface temperature (about the mean temperature) relative to the driving air-water temperature difference that existed over the course of the run. These images can be intercompared.

Mean thermal imagery is shown in figures 6-33 through 6-36. The results from the six grid A runs are presented, as are five of the six grid B runs that correspond to the same frequency-stroke combinations. Additionally, table 6.5 reports the environmental conditions for each run presented in these figures. Surface pressures were estimated in a manner described in section 7.3.2. Also shown in this group of figures are the respective standard deviation images computed from each sequence of 900 images. The temperature field in these images also was scaled by the appropriate air-water temperature difference: $\sigma_T/(T_w - T_a)$. The standard deviation images should be reflective of actual variability and not apparent variability due to imager noise. This is because the noise threshold of the

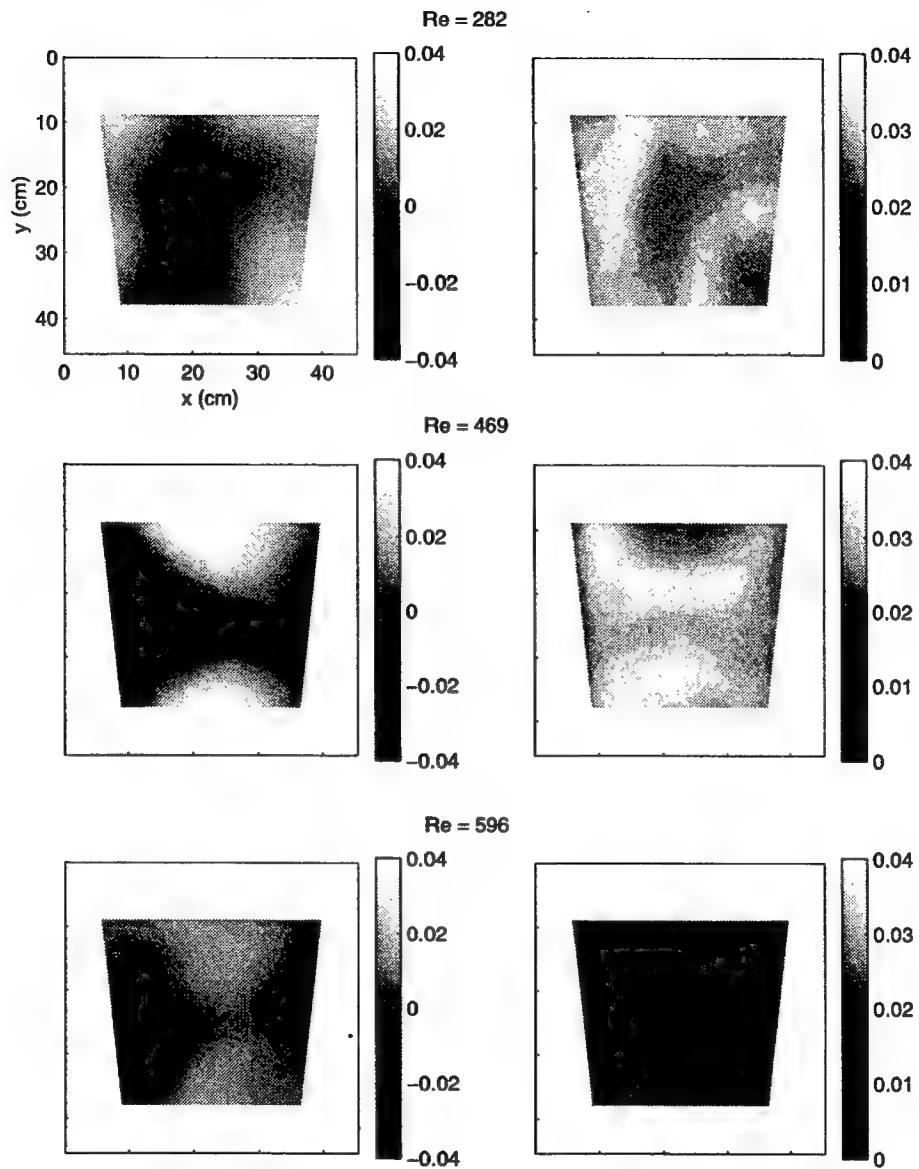


Figure 6-33: Persistent thermal features and their variability at the free surface for grid-stirred turbulence. Shown are results for $Re = 282, 469$, and 596 (grid A). The left panels are the mean thermal distributions. The contours are of dimensionless temperature deviation from the mean surface temperature, $(T - \bar{T})/(T_w - T_a)$. The right panels are the spatial distributions of standard deviation from the temporal mean. The contours are of dimensionless temperature variability, $\sigma_T/(T_w - T_a)$. The xy -plane is at the free surface and covers the complete plan section of the grid tank. The measurement domain is trapezoidal due to the angled, off-axis position of the IR imager.

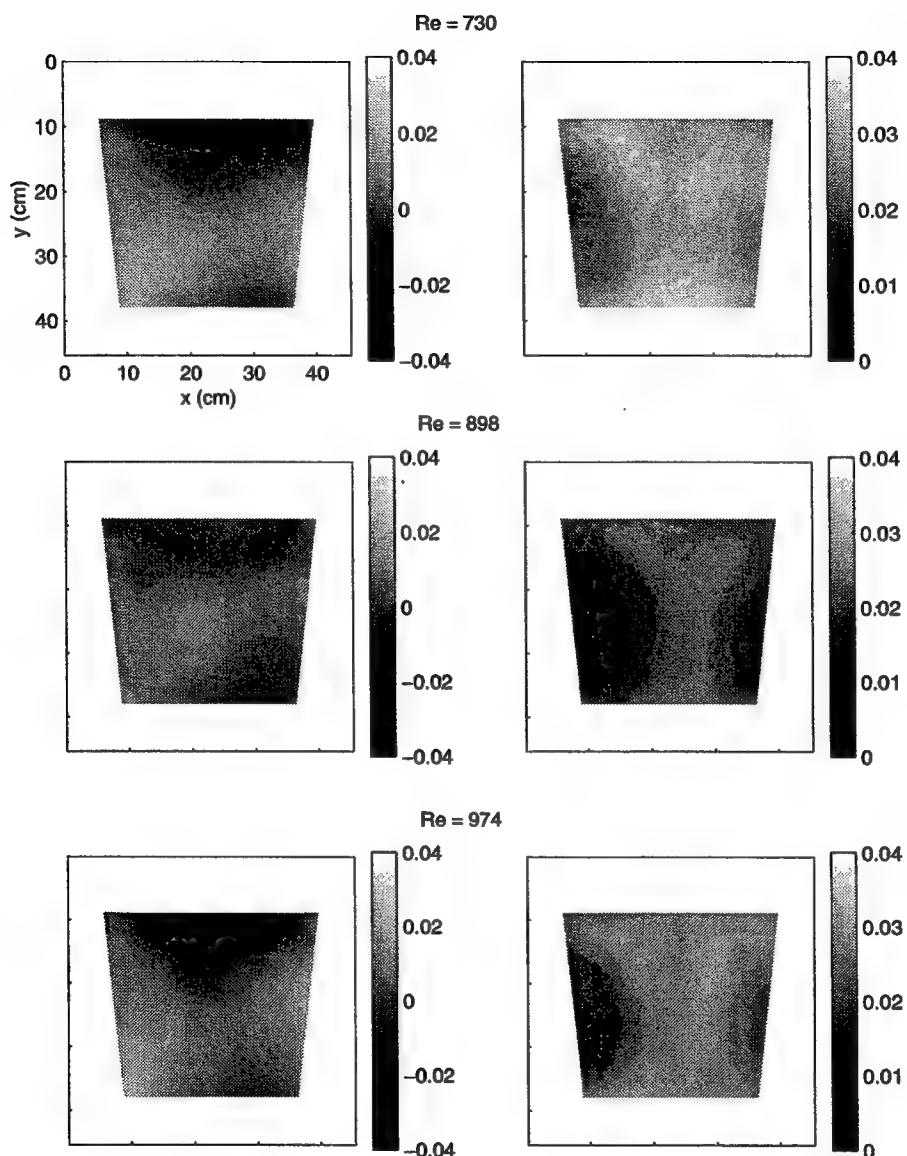


Figure 6-34: Persistent thermal features and their variability at the free surface for grid-stirred turbulence. Shown are results for $Re = 730$, 898 , and 974 (grid A). See figure 6-33 for details.

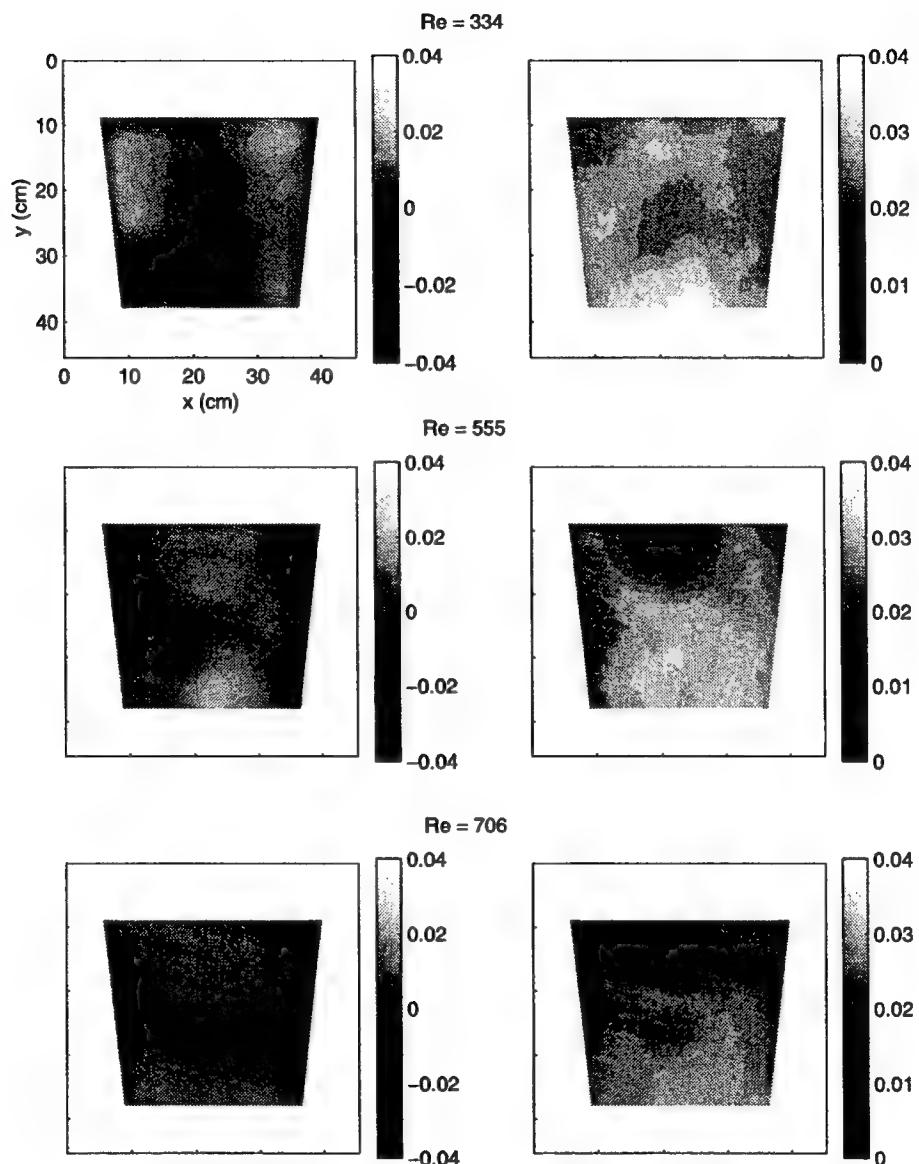


Figure 6-35: Persistent thermal features and their variability at the free surface for grid-stirred turbulence. Shown are results for $Re = 334, 555$, and 706 (grid B). See figure 6-33 for details.

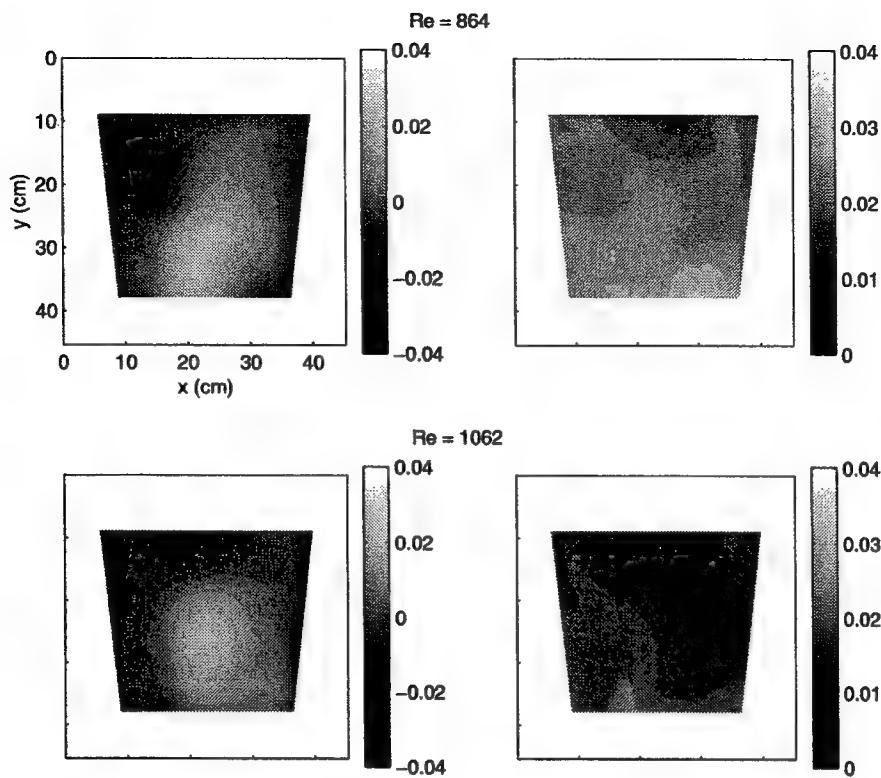


Figure 6-36: Persistent thermal features and their variability at the free surface for grid-stirred turbulence. Shown are results for $Re = 864$ and 1062 (grid B). See figure 6-33 for details.

imager was roughly 0.05°C , and the standard deviation values were typically on the order of $0.2\text{--}0.4^\circ\text{C}$.

Inspection of both the actual sequences and the mean images confirms the expectation that in some cases, the free-surface flow was not ideally homogeneous and persistent flow features existed at the surface. For the grid A runs, $Re = 282$ showed consistent upwelling bursts at the surface near the tank corner regions. This compares well with the ADV result at depth in figure 6-10. Both $Re = 469$ and $Re = 596$ showed two expansive upwelling regions—one for $y < 22\text{ cm}$, one for $y > 22\text{ cm}$ —existing in the tank. The $Re = 596$ case was more persistent, and this is reflected by the reduced standard deviation field. Reconciling the results for these two Reynolds numbers with that of the ADV is less obvious. Lack of agreement could be due to repeatability issues arising from the sensitivity of the mean flow

Re	T_a (°C)	T_w (°C)	π (mN/m)
282	22.40	35.90	0.01
469	23.55	40.40	0.18
596	22.45	35.00	0.38
730	23.55	38.30	0.19
898	22.50	33.90	0.35
974	23.25	37.20	0.41
334	22.43	33.50	0.09
555	24.00	36.80	0.92
706	22.85	35.40	1.20
864	22.70	34.50	0.74
1062	22.75	33.20	0.50

Table 6.5: Environmental parameters during the IR imagery experiments. The temperatures T_a and T_w refer to the mean air and bulk water values over the course of each run (30 minutes). The surface pressure π also reflects a mean estimate over each run and was computed from the measured surface tensions before and after each run. The 95% confidence intervals for all surface pressure measurements were ± 0.24 mN/m.

to initial conditions. Also, the two measurements are from two different horizontal planes, separated by 8 cm. Furthermore, the relationship between the surface thermal features and the near-surface bulk flow structure is not necessarily straightforward. The conditions at $Re = 730, 898$ and 974 all showed significant mean renewal over the entire imaged plane, and this is in qualitative agreement with the consistent upward mean velocities seen in figure 6-11. The mean thermal imagery for these cases shows modest deviations about the average temperature. This is because most individual images were dominated by complete renewal of the imaged interfacial area (most of the image was filled with warm fluid brought up from the bulk). Since this was consistent over an entire sequence, deviations from this behavior were small. This behavior, indicative of more uniform surface renewal, also was found in several of the grid B runs. The grid B condition $Re = 334$ showed distinct bursting regions similar to the corresponding grid A $Re = 282$ case. However, these bursting regions generally were not limited to the corner regions. A likewise comparison can be made between $Re = 469$ (grid A) and $Re = 555$ (grid B) in that the degree of renewal was similar but the spatial distribution of the renewal regions was not the same. For $Re = 706$, one or two large areas of significant renewal often were present; the locations of these renewal regions were

highly variable. The $Re = 864$ case showed slightly smaller and more numerous areas of significant renewal; the locations of these areas also were variable, with the central region of the imaged water surface being disrupted often. Lastly, the $Re = 1062$ condition showed expansive renewal across the entire imaged area. The imagery for $Re = 1152$ was distorted by ambient biasing effects that were not removed by the procedure discussed earlier, and is not shown.

Comparisons between the mean thermal imagery and the bulk PIV results for $Re = 596$ and $Re = 730$ are possible. As just stated, the IR imagery for $Re = 596$ showed two expansive upwelling regions, each occupying roughly half of the entire free surface. Recalling figure 6-21 and considering the turbulence field, it was found that the $y = 10$ cm and the $y = 35$ cm planes possessed greater turbulent energies than the midplane region. This disparity could explain the existence of the two regions of enhanced surface renewal. That observation, coupled with the almost universal mean upwelling in these two planes and downwelling in the center of the tank further explains the IR imagery. The results from the bulk PIV and the IR imaging for $Re = 730$ also seem to portray a similar flow field picture. The IR images indicate strong renewal across the imaged plane; the PIV results show fairly uniform turbulence in all three measurement planes and persistent upwelling throughout the tank except near the side walls. Neither measurement at this Reynolds number reveals a significant degree of spatial inhomogeneity. Instead, both sets of results show that the flow is reasonably uniform in space.

The standard deviation imagery in figures 6-33 through 6-36 is reflective of the temporal variability of temperature at the water surface. If temperature variability is assumed to be partially reflective of surface flow variability, some insight may be gained into the free-surface flows for the various forcing conditions examined. It is helpful again to consider what a result for an ideally homogeneous, random surface flow field would appear like. One would expect the standard deviation for the sequence to be uniform in space and obtain some finite value. The actual value would depend on the degree of variability of the flow at the surface, and also on the air–water temperature difference. The effect of the latter is essentially removed by the normalization described above. Thus, the standard deviation images should reveal how consistent the surface flows are. Consider as an illustrative example the case

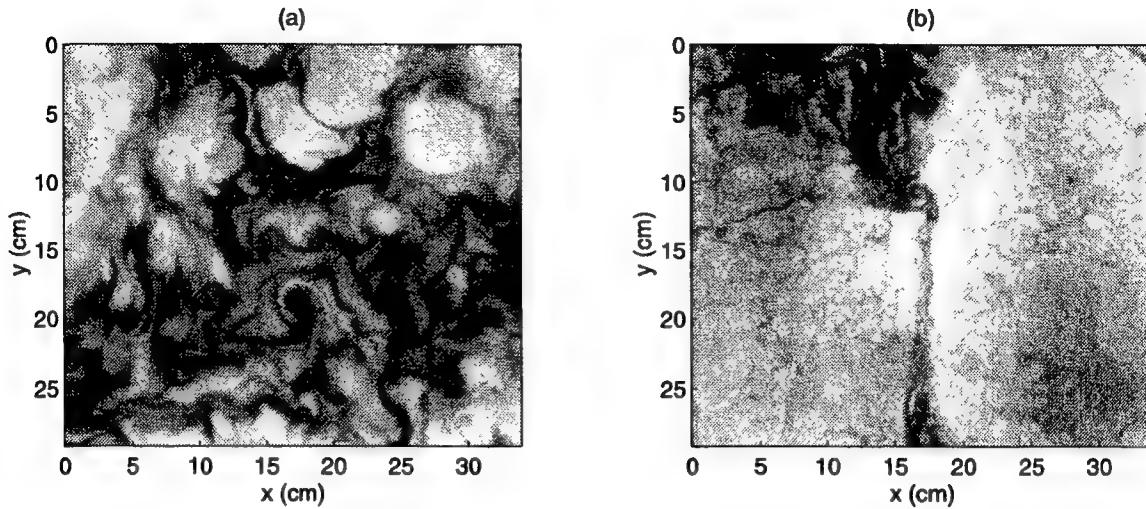


Figure 6-37: Comparison of surface renewal behaviors for two different grid conditions: (a) $Re = 334$ (grid B) and (b) $Re = 898$ (grid A). The axis origins are arbitrary.

for grid A, $Re = 282$. The deviation about the temporal mean in the central region, and even more so in a very localized region in the $x, y > 30$ cm corner, is the smallest. This would imply that these mean, persistent flows—upwelling in this corner, reduced motions in the center—are relatively constant. (This was confirmed through actual observations as well.) The other regions of the surface, while giving rise to a definite mean flow structure, are slightly more variable in time. Employing this interpretation, the spatial uniformity of the temporal variability of the flow improves with higher grid oscillation frequencies (e.g., $Re = 596, 898$). Additionally, the magnitude of the temporal variability is reduced with higher grid frequencies. This therefore would indicate that the flows for these cases were relatively consistent in both time and space. Comparing the results for the two grids, slightly better characteristics are exhibited by the grid B runs, as was found to be the case from the ADV measurements. However, the improvement is not necessarily significant.

To illustrate some of the rich activity at the free surface, as detected through infrared imagery, figure 6-37 shows two examples of very different surface renewal behavior. Figure 6-37a shows an image from the sequence for $Re = 334$ (grid B). In this case, the renewal is due to localized surface disruptions resulting from impinging eddy motions. The surface

renewal is therefore spatially sporadic and intermittent. In figure 6-37b, the case for an image from $Re = 898$ (grid A), virtually all of the imaged surface is undergoing renewal. The surface region is dominated by turbulent eddies disrupting the interface in an almost completely uniform fashion.

As a final point, it is necessary to discuss the role of free-convective effects on the observed surface thermal structure. Because of the existence of the cool skin layer, the cooler, denser fluid at the surface will tend to fall through the warmer, less dense fluid below, with warmer fluid then rising upward to replace the falling cooler fluid. When the temperature difference between the surface layer and the bulk fluid is significant (*e.g.*, large Rayleigh number), this can result in unsteady, turbulent falling sheets and plumes of cool fluid (Volino and Smith [134]). The contribution of free-convective motions relative to that of forced convection has been assessed previously (in the field) using a *surface* Richardson number (*e.g.*, Zappa *et al.* [146]). This surface Richardson number has been defined as

$$Ri_s = \frac{\alpha \nu g Q}{\rho c_p u_*^4}, \quad (6.14)$$

where α is the coefficient of thermal expansion for water, Q is the net air–water heat flux (positive into the water), c_p is the specific heat of water, and u_* is the waterside friction velocity (Kudryavtsev and Soloviev [75]). Equation 6.14 requires an estimate of the heat flux and the friction velocity: the former was obtained from the bulk water temperature time series and the expression, $Q = H \rho c_p (dT/dt)$, where H is the water depth; the latter was obtained using the velocity records from section 6.3.6 and the expression, $u_* = \sqrt{\langle uw \rangle^2 + \langle vw \rangle^2}$. In this way, the estimate of the heat flux is an upper bound, as some of the heat loss was likely through the tank walls and floor, in addition to through the water surface. The estimate of the friction velocity is also an upper bound since this measurement was at a depth of 8 cm, where the turbulence is slightly stronger than at the free surface (see figures 6-23 and 6-23). The surface Richardson number can be used to describe the ratio of buoyancy to shear forcing, and a critical value of Ri_s has been estimated for the case of field measurements to be on the order of -10^{-4} (*e.g.*, Soloviev and Schlüssel [123]; Zappa *et al.* [146]). Estimated values for the present data were in the range -1×10^{-4}

to -8×10^{-4} , depending on Re_{HT} and the air–water temperature difference. These values indicate that free-convective effects may play a role in the thermal structure at the water surface. It is likely that these effects are not insignificant for the lower Reynolds number cases but become less important for the higher Reynolds numbers. A more definitive assessment is not possible since the critical value noted above was estimated for wind-driven, open ocean conditions and may not extend to the situation for grid-stirred turbulence.

6.6 Summary

In this chapter, the grid tank designed, constructed, and utilized was described and documented. This tank is well suited for free-surface studies as particular attention has been paid to achieving a chemically clean system. The tank is also equipped for gas-transfer studies and can accept a number of sensors in a variety of locations. Diagnostics have shown that the grid flows are reasonably repeatable and quite stationary. However, the homogeneity of the flow is a potential concern. Measurements using all three technologies (ADV, PIV, and IR imaging) revealed bulk inhomogeneities as well as free-surface inhomogeneities. Before pointing out the implications of this finding, the issue of mean flows is discussed.

Historically, oscillating grid-stirred tanks have been described as “zero-mean-shear” or having no mean flow. This was one of the aspects that led to the selection of such a tank for this work. When secondary motions have been reported, typically, these features were not given significant attention. However, the results presented in this chapter have shown that these mean flows are a characteristic of grid-stirred tanks. It is felt that this particular tank is not in any way distinct in its behavior. To provide perspective, some earlier works using grid tanks are recounted, and the existence of non-negligible secondary motions is pointed out.

- Thompson and Turner [130] mention the presence of “steady flow patterns … in all parts of the tank.” It is found that, on average, their steady flow velocity scale was about 40% relative to their turbulent velocity scale (*i.e.*, $u/U \approx 2.5$). They also provide a comment that adds to the idea of flow sensitivity: “… the jet flows [from the individual grid bars] … combine in a pattern which is determined by the shape

of the tank, and by small differences in the strength of the flow from different parts of the grid.”

- In some of their initial experiments, Hopfinger and Toly [56] observed “strong secondary motions in the four corners.” Locating the mean grid position further from the tank floor was found to reduce these motions to “acceptable levels.”
- As already mentioned, McDougall [93] encountered noticeable mean secondary motions during his grid tank experiments. For his data, taken 10 cm from the grid for $f = 5$ Hz and $S = 1$ cm, it was found that $v/V = 3.4$ and $w/W = 2.3$ averaged over the available measurement points (20 for the horizontal component, 7 for the vertical). Additionally, McDougall points out inhomogeneities of the turbulence in planes parallel to the grid. The relative standard deviations $\sigma_v/\langle v \rangle$ and $\sigma_w/\langle w \rangle$ were 19% and 31% respectively, much like what was observed in the present experiments.
- In Kit *et al.* [71], the secondary mean flow was reported to be no more than 25% of the RMS turbulent velocity. This is not very different from what was shown for this tank in figure 6-15.

Interestingly, these accounts adequately describe the grid tank flows studied here, both qualitatively and quantitatively. Mean secondary circulations should therefore be considered an intrinsic feature of oscillating grid-stirred turbulence. Spatial non-uniformity is also inherent to these flows. However, both of these deficiencies can be relatively benign, provided they are recognized and sufficient spatial sampling exists to quantify the effects.

Because the free surface flow was found to be non-uniform in space, the plan for a single surface PIV measurement domain was abandoned. Instead, a number of surface locations were chosen for measurement. The locations were selected as representative of the various flow regimes found near, and at, the interface. There were two major implications of this modification: (1) most simply, more experimental trials needed to be performed, and (2) this provided an opportunity to examine the roles of different surface flow regimes on air–water gas transfer. The locations of the surface PIV measurements are presented in the next chapter.

Finally, the measurements in this chapter also explored the effects of different grid geometries on the grid tank flows. The results indicate that grid B ($M = 8.89$ cm) produces more spatially uniform flow with less dominant mean motions than does the grid A geometry ($M = 6.35$ cm). This is likely a consequence of the reduced solidity of grid B, which may limit the individual jet-jet interactions that can give rise to coalesced, larger mean jet-like flows (Hopfinger and Toly [56]). There is also evidence that grid A generates slightly higher levels of RMS turbulence for the same degree of forcing. This latter observation appears to be supported by measurements of higher gas-transfer rates for grid A over that for grid B at the same oscillation setting (see section 7.6). In light of these enhanced transfer rates, and with the possibility of using the surface flow inhomogeneity to study the effects of different flow regimes on interfacial transport, grid A was selected for the majority of the gas-transfer experiments in the next chapter.

Chapter 7

Free-surface hydrodynamics and air–water gas transfer

7.1 Introduction

Chapter 7 is the culmination of the efforts thus far discussed in this thesis. The surface PIV measurement technique is now coupled with measurements of gas exchange in the grid tank. Results and insights from the analysis of the grid tank flows in chapter 6 are incorporated. Free-surface hydrodynamics, air–water gas transfer, and surface films are all collectively studied in a series of experiments in the grid-stirred tank.

7.2 Earlier efforts and motivation for the present approach

The number of experimental efforts aimed at exploring scalar transport across air–water interfaces in seemingly endless. Some studies have focused solely on the waterside hydrodynamics or on the airside forcing. Others have looked exclusively at the interfacial transport or scalar concentration gradient near the interface. Some investigations have successfully examined both, revealing helpful relationships between the two. Inclusion (or acknowledgement) of surfactant effects often is sporadic, and rarely are the hydrodynamics, the transport, and surfactant effects treated simultaneously. Some of these efforts are surveyed next.

Experimental studies in open channel flumes have been conducted to explore the nature of the turbulent structures impinging at the free surface. Rashidi *et al.* [110] looked at non-wavy wind-sheared and shear-free channel flows with hydrogen bubble visualization, but did not make concurrent scalar transport measurements. Further work by the same group has continued to explore the structures and features for these wall-bounded flows, with a main thrust being qualitative descriptions of the various mechanisms involved (Rashidi [109]; Kumar *et al.* [77]). While these studies have provided useful conceptual illustrations of the turbulent interactions at a free surface, they are limited by the lack of supporting transport measurements.

An early, and ambitious, study by Broecker *et al.* [12] considered the various aspects of the interfacial transport problem with simple measurements of airside velocity profiles, waves, air- and waterside concentration profiles of CO₂, and surface film effects for a wind-wave tunnel flow. A result taken from this work is included as part of figure 7-1. The Broecker *et al.* result clearly shows the dramatic impact of surface films on gas exchange across air-water interfaces. It also shows the less than ideal performance of a wind speed parameterization of the gas-transfer velocity. More recent measurements under very controlled chemical conditions by Frew *et al.* [35] show both of these points very clearly. Also included in figure 7-1 is a result from the Frew *et al.* work showing the wind speed dependence and the effect of surface films on the transfer velocity. These results illustrate a current limitation of wind speed parameterizations. Without critical knowledge of the surface chemistry of the interface, the actual magnitude of the transfer velocity can be uncertain.

In addition to wind speed, the friction velocity is another often used parameter to characterize gas transfer. The friction velocity is arguably a more suitable parameter since it is associated with the processes near the air-water interface. Two results showing this parameterization are shown in figure 7-2. Similar to that for wind speed, these data also show noticeable scatter in the transfer velocity for fixed friction velocities. In figure 7-2a, the scatter is attributed to surface films. In figure 7-2b, the source of the scatter is less evident, since different tank geometries and fetch, for example, also have been shown to yield varying results for the transfer velocity (Jähne *et al.* [63]). While not perfect, wind speed and friction velocity are presently the most readily measurable quantities in the ocean

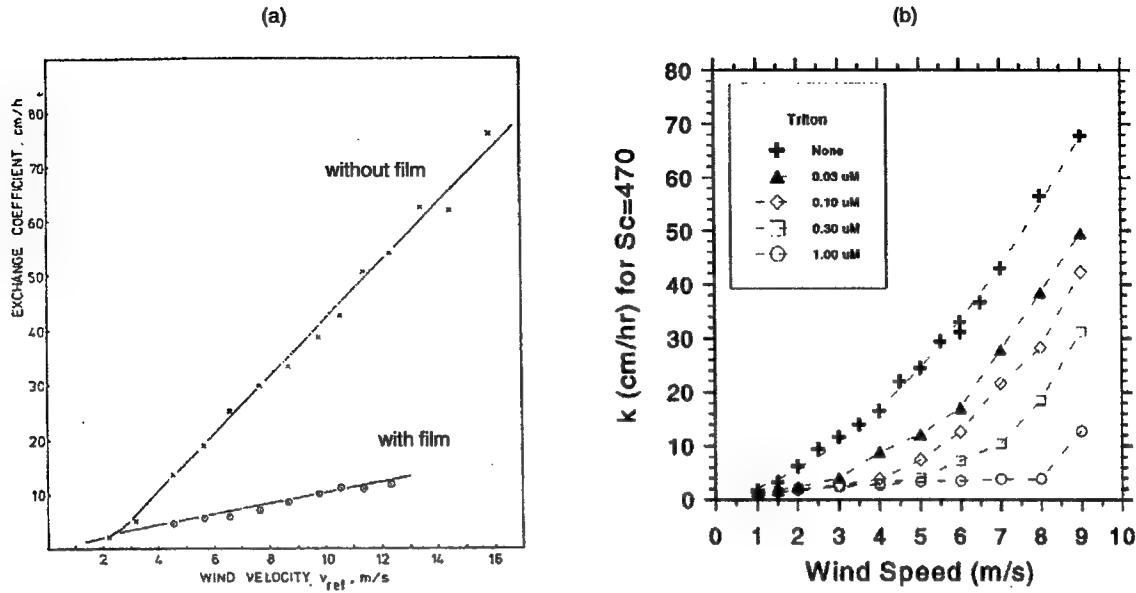


Figure 7-1: Gas-transfer results from (a) Broecker *et al.* [12] and (b) Frew *et al.* [35] showing wind speed dependence and surface film effects. The data in (a) were collected in an 18 m wind-wave tunnel. The exchange coefficient is equivalent to the gas-transfer velocity, k , in this case for CO_2 , and the film used was oleyl alcohol. The data in (b) were collected in an annular wind-wave flume (0.5 m nominal diameter). The figure key shows the different dilute solutions of the soluble surfactant Triton X-100. The transfer velocities shown are for O_2 ($Sc = 470$).

environment.

The wind-wave tunnel experiments by Komori *et al.* [73] measuring CO_2 reaeration and air- and waterside velocities have shown some success in applying a surface renewal model to the transport process by estimating the frequency of appearance of organized motions near the interface. A somewhat unique study in a moving bed flume by Gulliver and Tamburrino [127] did not look at transport, but did attempt to measure the wavenumber spectra of the surface divergence. The flow measurement technique (see section 4.1) was used in a manner that imaged a large area (appearing to be about 70 cm square), which leads to questions about spatial resolution. It is therefore difficult to interpret the significance of the results. However, the approach taken by Gulliver and Tamburrino is a thoughtful one, owing to the expected importance of surface divergence in interfacial transport.

Experiments in small tanks with jets or oscillating grids have approached the problem

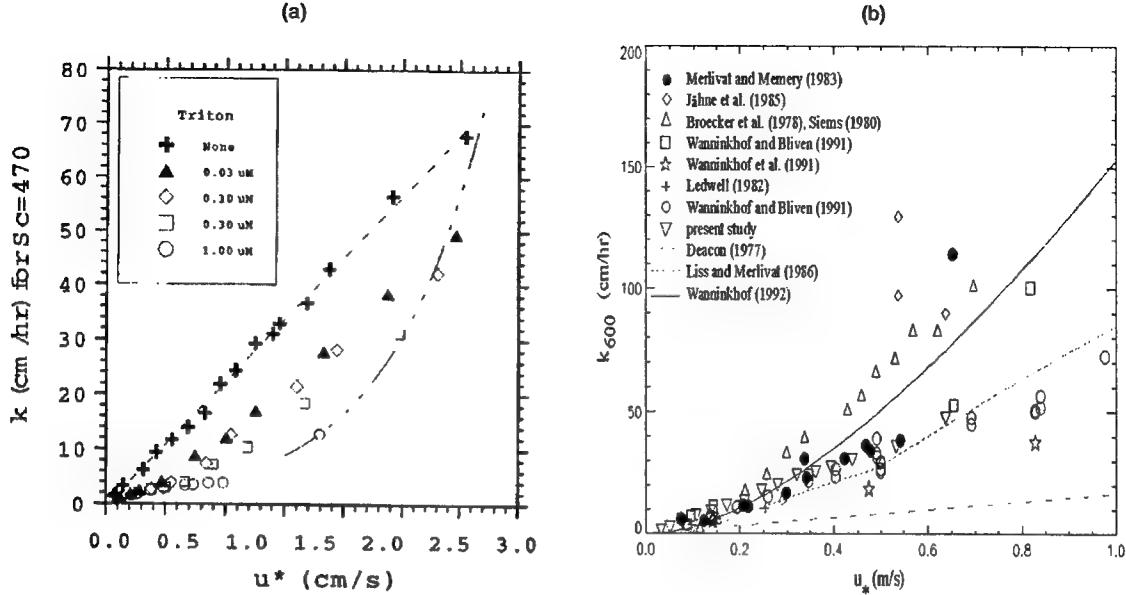


Figure 7-2: Gas-transfer results from (a) Frew *et al.* [35] and (b) McGillis [95] as a function of friction velocity. The data in (a) are as in figure 7-1b and u_* is the waterside friction velocity. The data in (b) were collected from numerous laboratory studies; the transfer velocities have been normalized to that of CO_2 ($Sc = 600$).

from the consideration of the waterside hydrodynamics. George *et al.* [41] have studied the problem of gas transfer in a vessel agitated by an array of microjets with and without a surface film. Using LDV measurements to quantify the turbulent energy levels near the interface, they show a relationship between the turbulent energy and the transfer velocity. However, the presence of a surfactant gives rise to variability in the transfer velocity for a given level of “interfacial” kinetic energy. Khoo and Sonin [70] looked at a similar problem in a vessel with a single upflowing jet. Their results for CO_2 exchange indicate a two-regime transport process. At low turbulent intensities, one hydrodynamic relationship is found to describe the data, while for higher intensities above a certain “break point,” a second relationship is needed to account for the data. The authors fail to explain the physical significance of the break point, which is most likely just the manifestation of adventitious surface film effects. Such behavior has also often been seen in wind-wave studies (*e.g.*, figure 7-1b, where many of the surfactant curves show a noticeable kink, indicating a transition; also, figure 7-3).

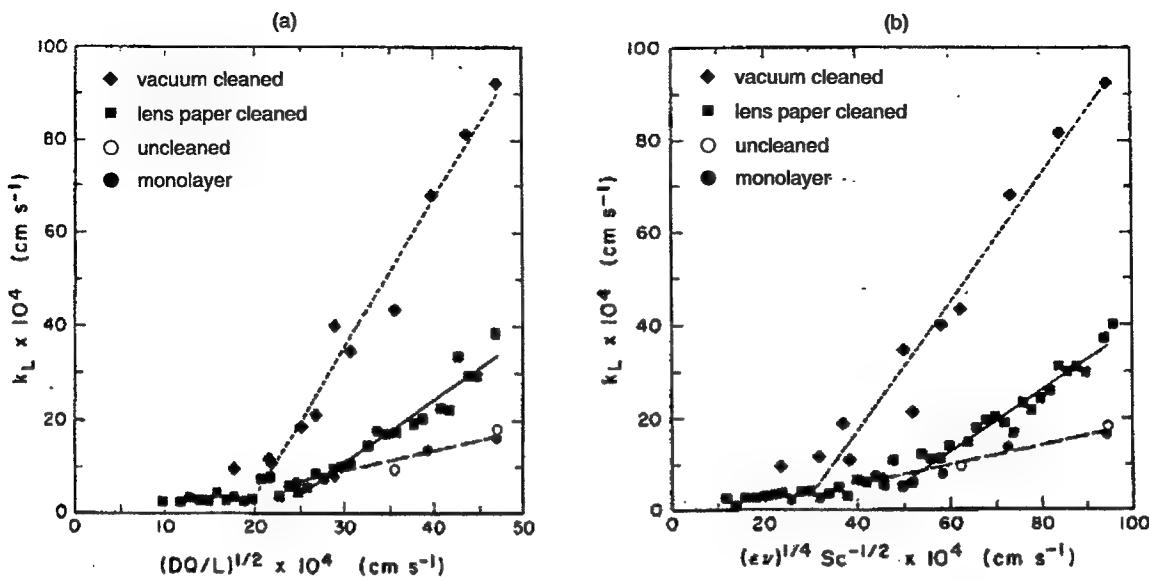


Figure 7-3: Gas-transfer results from Asher and Pankow [6] as a function of turbulent eddy parameters. (a) Gas-transfer data plotted against large eddy model of Fortescue and Pearson [33] (equation 2.51). (b) Gas-transfer data plotted against small eddy model of Lamont and Scott [78] (equation 2.52). The monolayer was of 1-octadecanol.

Dickey *et al.* [28] made use of a streak photography method to quantify the waterside turbulence while simultaneously measuring gas-transfer rates for a number of gases in an oscillating grid-stirred tank. Although innovative, the results would now be considered somewhat crude as the spatial resolution of the velocity measurements was poor and measurements were not made very near the interface. Chu and Jirka [18] made concurrent measurements of near-surface profiles of turbulent velocities and O₂ concentration, also for grid-stirred turbulence. Likewise, Asher and Pankow [6] used a laser induced fluorescence (LIF) technique to obtain bulk CO₂ concentration measurements for a grid-stirred system with and without a surfactant, and for various surface cleaning procedures. They did not make measurements of the turbulence, but instead relied on the empirical relations of Hopfinger and Toly (see section 6.2.3), which, as already seen, are not necessarily accurate near the air-water interface (section 6.4). The inaccuracy becomes worse for surfactant-adsorbed surfaces (*e.g.*, see figure 7-13). Results from the Asher and Pankow work are shown in figure 7-3. In figure 7-3a, the transfer velocity is plotted against the predic-

tion of the large eddy parameterization of Fortescue and Pearson [33]; in figure 7-3b, the transfer velocity is plotted against the small eddy parameterization prediction of Lamont and Scott [78]. In both cases, the transfer velocity is widely scattered for a fixed hydrodynamic parameter. Even though the values for Q , L , and ϵ that appear in the models (see section 2.4.3) were evaluated at the interface, the relations used are not necessarily appropriate in the interfacial region, especially when a surface film is present. Hence, the transfer velocity is not uniquely predicted by such parameterizations, much like the results with microjets of George *et al.* mentioned above. The debate over the controlling turbulent scales, small versus large, continues. It is most likely that neither model can individually explain the true physical situation sufficiently, but rather a (conditional) combination of the two is required (e.g., the “chain saw model” described by Moog and Jirka [102]). A final note on the Asher and Pankow results regards the cleaning methods that were explored. As is clear from the figure, the “vacuum cleaning” produced the highest transfer velocities, or by inference, the cleanest surfaces. This is the same surface cleaning method employed for the grid tank experiments in this thesis.

A good deal of effort has gone into exploring the role of waves on gas transfer. From some of the more recent efforts, wave slope has emerged as a relevant parameter. Wave slope has been measured using optical refraction methods like those described by Lange *et al.* [79]. Three-dimensional (two spatial dimensions and time) wave slope has been successfully measured with a field-deployable scanning laser slope gauge by Bock and Hara [11] and recently by Suoja [126] using a resistive wave-wire array buoy. Wave slope parameterizations of the transfer velocity by Hara *et al.* [48] and Jähne *et al.* [63] have shown favorable results for varying surface conditions. The good correlation between the two is likely rooted in the mutual sensitivity of wave amplitude (damping) and surface renewal processes to surface film effects. However, wave slope is not a panacea and there are limitations associated with it as well. Figure 7-4 shows two results indicating potential complications with wave slope. In figure 7-4a, the k -slope relationship appears to exhibit a tank geometry dependence, and in figure 7-4b, the dependence on wave frequency is evident.

A novel approach to the gas-transfer problem has been the use of heat as a proxy tracer for mass. This approach has been pioneered by Haussecker *et al.* [50] and Jessup *et*

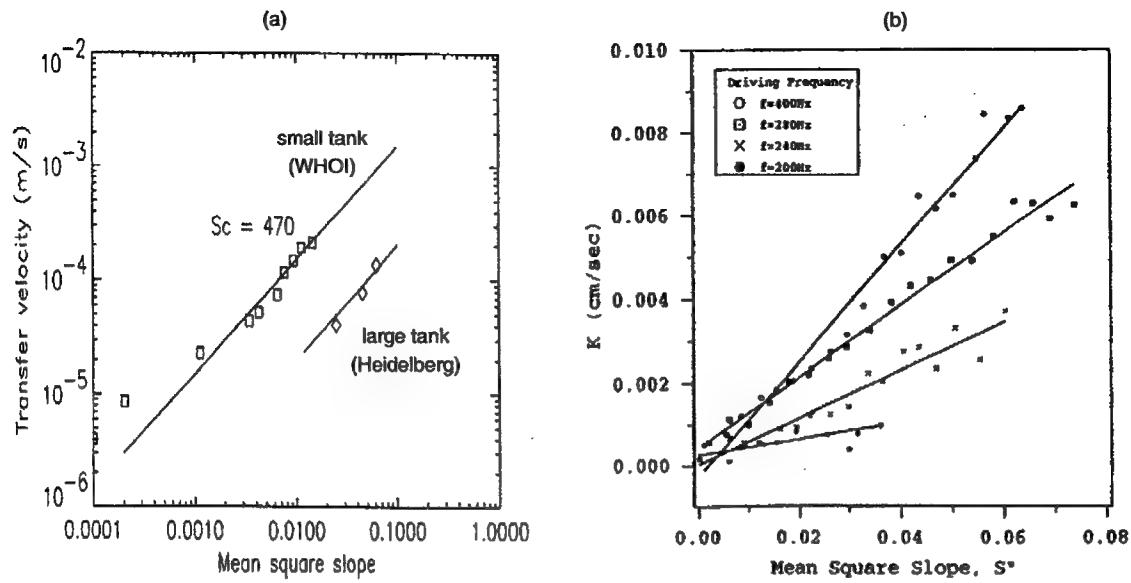


Figure 7-4: Gas-transfer results from (a) Hara *et al.* [48] and (b) Saylor and Handler [115] as a function of mean square slope.

al. [66] (also, Zappa [145]) with passive infrared (IR) surface imaging techniques and the related controlled flux technique (CFT). In the former approach, IR imagery of the water surface is used to reveal localized areas of surface renewal via disruption on the interface by microscale breaking events. Good correlations have been found between the fractional area affected by this surface renewal and the gas-transfer velocity with and without surface films. CFT techniques probe gas transfer by observing the response of the water surface temperature, and its spatial variability, due to an imposed heat flux (from a laser source). The CFT method currently relies on statistically-based surface renewal models to provide closure. The CFT approach is a significant improvement in the temporal resolution of gas flux measurements—time scales for this type of measurement are on the order of seconds. The concept of heat as a proxy for gases is currently an active area of research and certain details of the measurements are current topics of investigation.

At present, researchers are seeking a measurable quantity (that can be measured without prohibitive difficulty, particularly at sea) that accurately estimates the transfer velocity of scalars of interest (*e.g.*, climatically important gases, heat, water vapor, and pollutants).

As already seen, it is imperative that such a quantity properly capture the effects of surface films. This requirement has been a major reason for the current shortcomings of the present methods of parameterization. Because the influence of surface films is highly localized at the air–water interface, measurements away from this thin region will be continuously susceptible to certain limitations. Wave slope has emerged as one of the more robust parameters, although there is some uncertainty as to which waves should be considered—the full measured spectrum, the low wavenumber spectrum, the high wavenumber spectrum, or some combination. The fact that wave slope is a geometric characteristic of the surface, rather than directly related to the physical mechanism of transport is also a potential drawback. A more fundamental mechanism like microscale breaking is therefore a more attractive alternative and is growing in favor. Very recent advances in airside measurement techniques, such as the eddy correlation technique, the gradient technique, and the relaxed eddy accumulation technique, have yielded encouraging results for temporally resolved ($\mathcal{O}(15$ minutes)) measurements of air–sea gas fluxes (Asher and Wanninkhof [5]). However, some of these methods are not free from exciting debate.

This chapter aims at further understanding the critical hydrodynamic processes involved in air–water mass transport. The works of Zappa [145] and Gulliver and Tamburrino [46] are most closely related to the present effort. In those studies, and in this thesis, the *physical mechanisms* behind air–water transport are explored. In the present work, measurements of the fluid velocity field at the free surface are made in order to reveal surface renewal events quantitatively. Concurrent gas-transfer measurements are made, and the role of surface films is considered closely.

7.3 Characteristics of the water surfaces studied

In this section, results from laboratory analyses with representative samples of the water surfaces used in the gas-transfer experiments discussed subsequently are presented. Laboratory tap water and commercial spring water are compared. The impact of PIV particles is also assessed, and the surfactants studied are examined.

7.3.1 Water selection

As mentioned in section 6.2.4, commercially available spring water was used for the grid tank experiments. This choice was based on chemical properties and practicality considerations. Tap water was ruled out instantly as it is well known to contain a variety of unwanted organic material. Tap water can also show variability over the course of a day and typically shows seasonal dependence. Therefore, tap water should never be used in situations where surface cleanliness is a concern. Distilled water often contains significant levels of surface-active organics that can adsorb at the water surface. The presence of these adsorbing contaminants is reflected by the time-dependent behavior of the gas transfer. Skimming and/or vacuuming techniques may remove these surfactants. However, surface films can quickly redevelop when these procedures are ceased. Filtered deionized water has also been found to exhibit similar surfactant adsorption behavior. Laboratory analysis has shown that some commercially available spring waters, while not being surfactant free, contain very low amounts of organics and can be significantly cleaner than laboratory distilled water (Frew *et al.* [35]; DeGrandpre *et al.* [27]). The particular spring water used was chosen because it was available in large bulk quantity.

7.3.2 Surface tension and surface pressure measurements

As a starting point, and for use as a reference for surface pressures, measurements of the surface tension of the spring water as a function of temperature were collected. These were obtained using the Wilhelmy plate method as described in section 2.3.1. Spring water from a single jug was dispensed into a clean beaker and either chilled in a refrigerator or heated using a hot plate to near the desired temperatures. The results are shown in figure 7-5. Also shown in the figure is the ultra-high pure (UHP) water result given earlier in section 2.3.1. It appears that the spring water surface tension is roughly 0.5 mN/m less than that expected for UHP water over this temperature range and under these conditions. The least-squares fit to the data is used throughout this thesis as the reference surface tension for the spring water:

$$\sigma_o(T) = 75.558 \frac{\text{mN}}{\text{m}} - 0.153 \frac{\text{mN}}{\text{m}^\circ\text{C}} T, \quad (7.1)$$

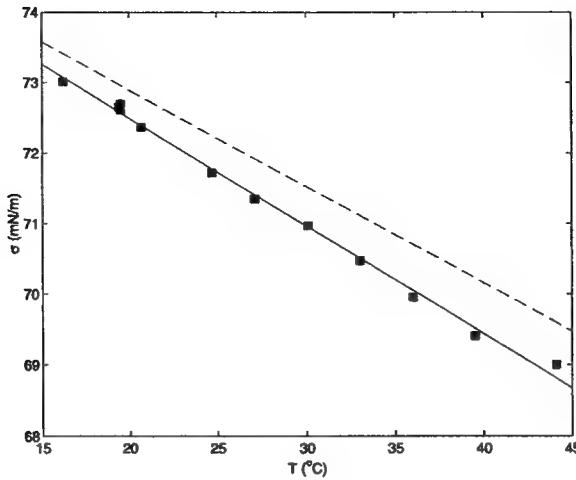


Figure 7-5: Surface tension of spring water as a function of temperature. The squares represent the measured values, the solid line represents the least-squares fit, and the dashed line represents the ultra-high pure water result from section 2.3.1.

with T in degrees Celsius.

In order to gain insight into the surface activity of the water surfaces explored during this thesis, surface pressure-area (π - A) isotherms were determined for a number of cases. The π - A curves were generated using a Langmuir-Blodgett film apparatus described in Goldman *et al.* [42]. The instrument was calibrated initially with UHP water. A repeated calibration measurement indicated an accuracy on the order of ± 0.10 mN/m. In the work of Goldman *et al.*, the instrument accuracy was estimated as ± 0.16 mN/m (2 standard deviations). In all cases, 500 ml samples were introduced into the pre-cleaned Teflon instrument trough. The initial and final trough areas were 714 cm^2 and 42 cm^2 , respectively, and the film compression rate was 50 mm/min.

To compare laboratory tap water and the spring water, two π - A curves were determined for each. For one set of experiments, each water sample was introduced into the trough and allowed to stand for 20 minutes (“20 minute film”). For the second set of experiments, each sample was allowed to stand for 2 hours (“2 hour film”). The results from these tests are shown in figure 7-6. Both the 20 minute film results and the 2 hour film results illustrate the surface-active nature of common tap water. The constituents of the tap water appear to make it behave as an expanded film (see below). Time-dependent behavior is found for both

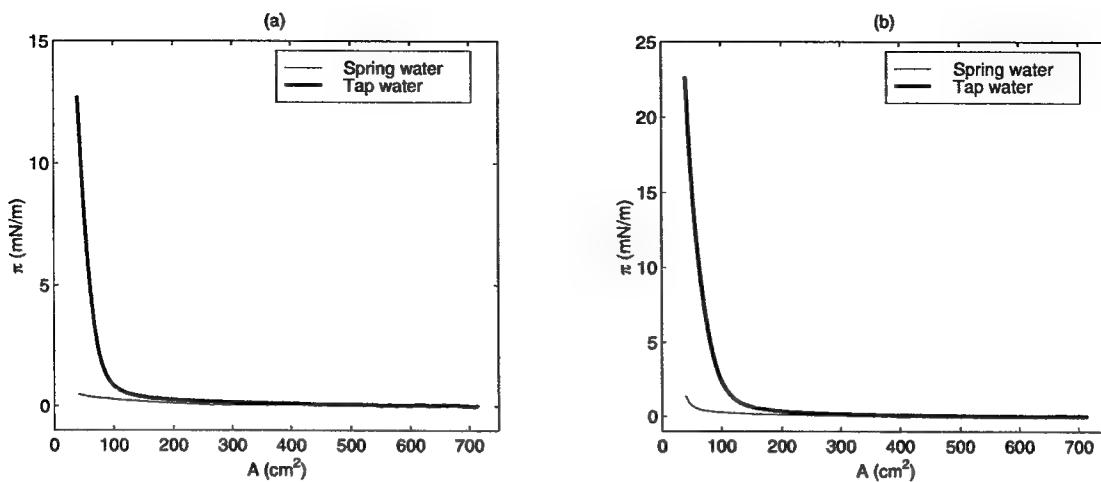


Figure 7-6: Surface pressure isotherms for laboratory tap water and commercial spring water. (a) 20 minute film results, (b) 2 hour film results. Note, the scales for surface pressure π are different between (a) and (b). The area, A , corresponds to the film area.

water samples, evidenced by the increase in surface pressure from 20 minutes to 2 hours. The final 2 hour surface pressure for the spring water was only 1.38 mN/m. Because it is quite significant, the time-dependent behavior for the tap water sample can be assumed to be dominated by surface adsorption of bulk surface-active material rather than by possible airborne contamination (Frew [39]).

The two surfactants used in the gas-transfer experiments were cholesterol (5α -cholesta- 3β -ol) and oleyl alcohol (cis-9-octadecan-1-ol, $C_{18}H_{36}O$). Both of these surfactants are insoluble in water and form a surface film when introduced to the water surface. The choice of insoluble surfactant over soluble was made in order to focus on surface phenomena and exclude possible bulk fluid effects. Depending on the chemical nature of the surfactant molecule, a monolayer may be classified into three types: (1) solid, or condensed, (2) liquid, or expanded, or (3) gaseous (Gaines [40]). Cholesterol is a solid-like film, oleyl alcohol is a liquid-like film. Inadvertently, over the course of gas-transfer work in the grid tank, the cholesterol sample being used became contaminated to an unknown degree. As a result, the films spread using the sample cannot be considered strictly cholesterol, and therefore these results are called “cholesterol,” in quotes. For the surfactant tests, 50 μ g samples were diluted in hexane for effective film spreading. In an effort to mimic grid tank conditions, the

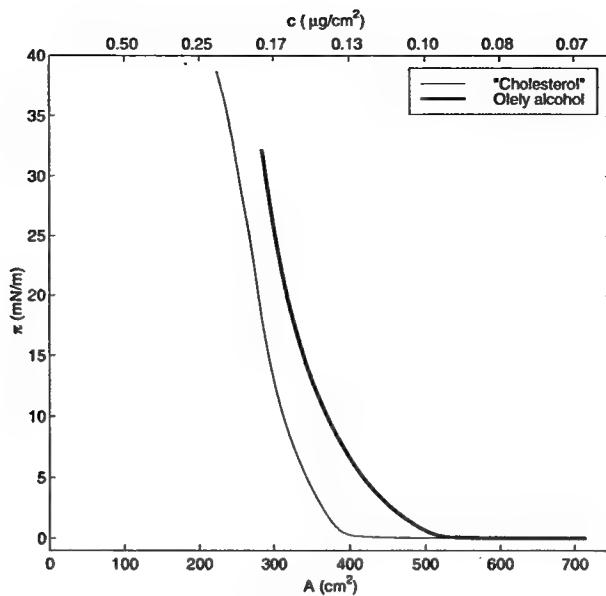


Figure 7-7: Surface pressure isotherms for films of “cholesterol” and oleyl alcohol. The data are for 20 minute films using 50 μg samples of each. The upper abscissa reflects the film surface concentration.

trough was filled first with 500 ml of spring water and each sample was spread incrementally onto the water surface using a micro-pipette. The hexane solvent subsequently evaporated in a matter of minutes. For these trials, only 20 minute films were tested. Figure 7-7 shows the surface pressure results. The oleyl alcohol isotherm is as expected. However, the “cholesterol” isotherm does not reflect a solid-type film result, showing more liquid-like behavior and confirming the existence of a contaminated sample. Both films completely collapse near film areas of 200–300 cm^2 . Upon collapse, the film buckles and the surfactant molecules are driven from the surface to form agglomerates of an adjacent bulk phase (Gaines [40]).

The final set of surface pressure isotherms collected were for spring water with PIV particle seeding (see section 4.2.2). Tests were conducted with PIV particles that had undergone a chemical cleaning procedure because of early evidence that the raw particles, which are packaged as a concentrated water mixture, were a source of surface-active material. This evidence was that grid tank gas-transfer measurements showed a marked reduction in the flux of O_2 for a given grid forcing when PIV particles were in use (see figure 7-24). By them-

selves, the particles should not have been a significant source of contamination. Acrylic is reasonably inert and the dyes used for fluorescence are relatively benign as well. However, the specific details of the manufacturing process and handling were unknown, and possible contaminating residues could be present. Additionally, the water that the particles were shipped in could not be assumed clean.

The cleaning procedure for the PIV particles was as follows. Raw PIV particle solution was poured incrementally into a filtering apparatus that had a sintered glass frit (nominal pore size greater than $5 \mu\text{m}$) at the bottom. A 50% methanol, 50% UHP water solution then was flushed through the filtering apparatus. Additional PIV particle mixture was added, and more of the 50:50 methanol/water solution was flushed through. This procedure was repeated until all the PIV particles had been filtered, and the entire process was repeated two additional times. After the three filtering operations, the PIV particles were removed from the apparatus and put into a final mixture with UHP water.

For the surface pressure isotherm tests, a sample containing actual PIV particles could not be used since the particles were found to wet the Wilhelmy plate and collect along its perimeter, confusing the measurement. In lieu of the actual mixture, only the supernatant of the cleaned PIV mixture was tested. To do this, the cleaned PIV particle mixture was filtered to separate the particles from the supernatant. A mixture of cleaned PIV supernatant and spring water was created in an attempt to represent the grid tank situation. The seeding in the grid tank PIV experiments was never highly systematic. Seeding and image quality were monitored subjectively, as is often done in PIV experiments. Therefore, replicating the exact seeding density conditions for the grid tank work was not possible. However, a crude estimate, is that 500 ml of PIV seed particle mixture was added to the tank during experiments. Using this amount and the tank water volume, it was determined that 2.5 ml of the PIV supernatant should be added to 500 ml of spring water, and the result introduced into the trough for analysis. The effect of the cleaned PIV particles on the surface activity of the spring water is illustrated in figure 7-8. The final surface pressure for the cleaned PIV supernatant mixture was found to be 0.35 mN/m higher than the final pressure for the pure spring water. The effect of the presence of the cleaned PIV particles is therefore noticeable, yet slight. Because the details of PIV particles may often be unknown,

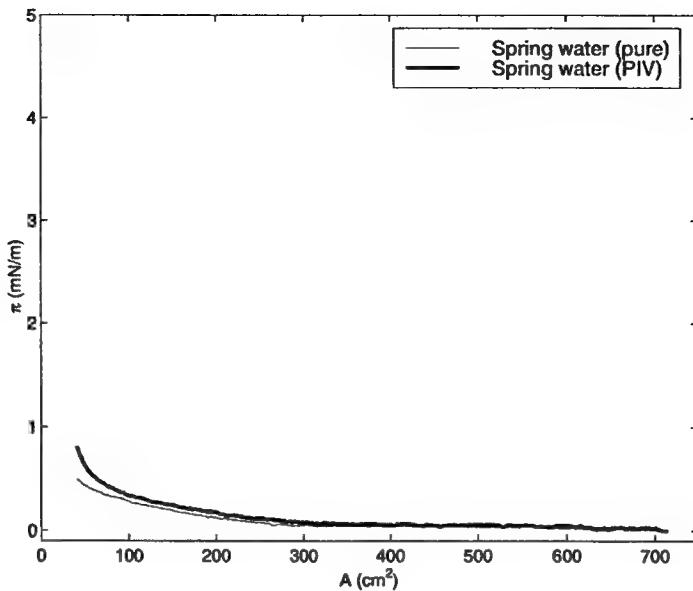


Figure 7-8: Surface pressure isotherms for pure spring water and spring water with cleaned PIV supernatant. The data are for 20 minute films. See text for further details.

as in this case, it would be considered prudent to always consider cleaning PIV particles for free-surface measurements where interfacial cleanliness is a concern.

As a final measure, statically derived surface elasticities were compared. The static elasticities were computed directly from the π - A data using equation 2.35. Only the 20 minute film results were considered. For the pure spring water and the cleaned PIV supernatant in spring water, the peak elasticity results were very low (less than 1.8 mN/m at maximum film compression). The two surfactants and the tap water sample showed significant elastic behaviors. The results for these three cases are given in figure 7-9. The oleyl alcohol sample begins to show noticeable elastic behavior at a film area of about 530 cm^2 ($0.094 \text{ } \mu\text{g}/\text{cm}^2$). At the point of film buckling, the elasticity is near 120 mN/m. The “cholesterol” sample starts exhibiting elastic behavior near a film area of 410 cm^2 ($0.122 \text{ } \mu\text{g}/\text{cm}^2$). A peak elasticity occurs near 280 cm^2 ($0.179 \text{ } \mu\text{g}/\text{cm}^2$), where the value is 116 mN/m. For greater film compression, the elasticity falls off, most likely indicating surface-active molecules being forced into the bulk. The tap water does not show obvious elastic properties until it reaches a film area of roughly 100 cm^2 . The maximum elasticity measured for the tap water was

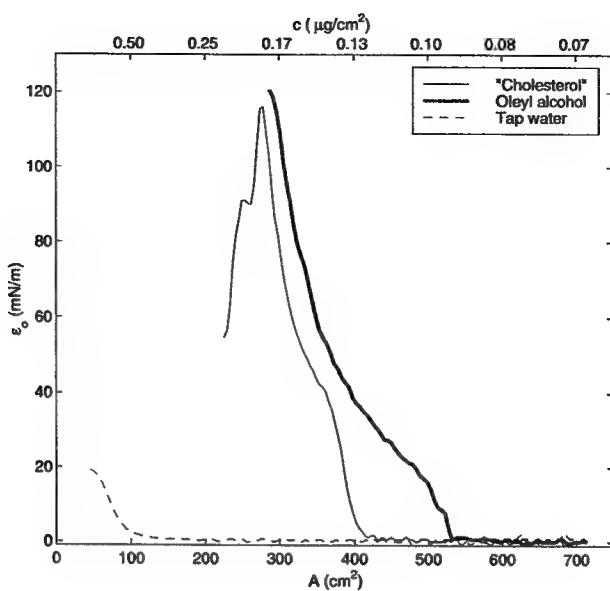


Figure 7-9: Statically derived elasticity isotherms for the two surfactants studied (“cholesterol” and oleyl alcohol) and laboratory tap water. The upper abscissa reflects the estimated surface concentration for the two surfactants.

close to 19 mN/m.

It is quite obvious that the addition of a surfactant to a water surface endows that surface with certain new rheological properties. These properties often give rise to markedly different surface flow behavior as mentioned in chapter 2 and can have a profound impact on air–water exchange rates. However, proper detection of the presence of such surface films is difficult, particularly when the setting is not a controlled chemistry laboratory. To illustrate this point, consider the initial surface pressures (*relative to UHP water*) for the various surfaces studied in the Langmuir-Blodgett film apparatus (see table 7.1). These surface pressures correspond to the surface tension measured after initially waiting 20 minutes and before film compression. The UHP water result in table 7.1 is from a measurement with UHP water after the trough apparatus had been calibrated with another sample of UHP water. Therefore, the uncertainty of the surface pressures shown in table 7.1 is on the order of 0.1 mN/m. Based on these numbers, the spring water, tap water, and the cleaned PIV supernatant solution all appear equally “clean” within the measurement uncertainty. Comparing the spring water and the surfactant films, the surface pressure difference is

Surface	π_i (mN/m)
UHP water	0.10
Spring water	0.39
Tap water	0.35
Spring water (PIV, cleaned)	0.28
“Cholesterol”	0.58
Oleyl alcohol	0.41

Table 7.1: Initial surface pressure π_i relative to UHP water for various surfaces studied. Equation 2.31 was used for the reference surface tension (σ_0). The surface pressures correspond to the surface tension measured after initially waiting 20 minutes and before film compression.

only a few tenths of a mN/m. The point these data make is that detection of the *slight* presence of a surface-active material at a free surface requires extremely sensitive measurements, and even with such measurements, distinguishing between degrees of cleanliness is potentially ambiguous. This leads to the confounding consequence that while two water surfaces may appear equal in terms of surface pressure, one may give rise to significantly different gas transfer compared to the other. This difficulty was experienced with the grid tank gas-transfer measurements in this thesis. Contributing to this difficulty was the slightly greater measurement uncertainty associated with the surface tensiometer used in the grid tank experiments (nominally ± 0.28 mN/m). This tensiometer utilized a simple mechanical weighing balance that is not as accurate as more sophisticated and expensive electro-balances currently available.

7.3.3 Surface properties of natural water surfaces

The water surfaces studied in this thesis were highly artificial when compared to actual water surfaces in nature. The surfaces of lakes and seas often are made up of a complex mixture of surface-active compounds, both soluble and insoluble. The results of Frew and Nelson [38] have shown that there exists considerable variability among surface films from different slicks and films derived from subsurface samples. The response of these mixed natural films is generally *viscoelastic*, and the π - A behavior is commonly not reversible upon repeated compression and expansion. The result is a film-influenced surface that

exhibits a complex stress response to compressional and dilatational straining.

Studies looking at the surface physical properties of sea-surface films have indicated that these films generally behave as expanded liquid or gaseous type films (Frew [34]). Figure 7-10 includes sample π - A curves for a number of sea-surface films adapted from Frew and Nelson [37]. As mentioned, oleyl alcohol represents a liquid film, and the “cholesterol” film was also found to exhibit liquid-like behavior. Static elasticities for films collected at coastal and open ocean sites have yielded maximum values in the range of 20–40 mN/m for surface pressures in the range of 10–15 mN/m. For surface pressures of 0.5–1.0 mN/m, elasticities of 2–8 mN/m have been reported (Frew [34]). Frew comments that surface pressures typical for the ocean surface are quite small: $\pi < 1$ mN/m. Values of surface pressure measured in visible ocean slicks using the spreading drop method at coastal and open ocean sites are generally on the order of 1 mN/m (Hunter and Liss [62]). Based on the π - A results of the previous section, oleyl alcohol and “cholesterol” could be expected to exhibit slightly stronger elastic behaviors than natural waters. Therefore, one might anticipate that these types of insoluble surfactant would have a greater impact on free-surface hydrodynamics and interfacial gas exchange than would soluble surfactants. However, there is evidence that soluble surfactants can be *more* effective in retarding gas exchange under turbulent conditions (Springer and Pickford [125]). Enhanced recovery times and the more rapid reestablishment of soluble films after disruption by turbulence has been suggested as an explanation.

7.4 Grid-stirred free-surface turbulence

This section presents the results from the surface PIV measurements in the grid tank. The experimental approach and data analysis procedures are outlined, the data sets that will be shown in this chapter are introduced, the free-surface hydrodynamics are explored, and the effects of a surfactant are examined.

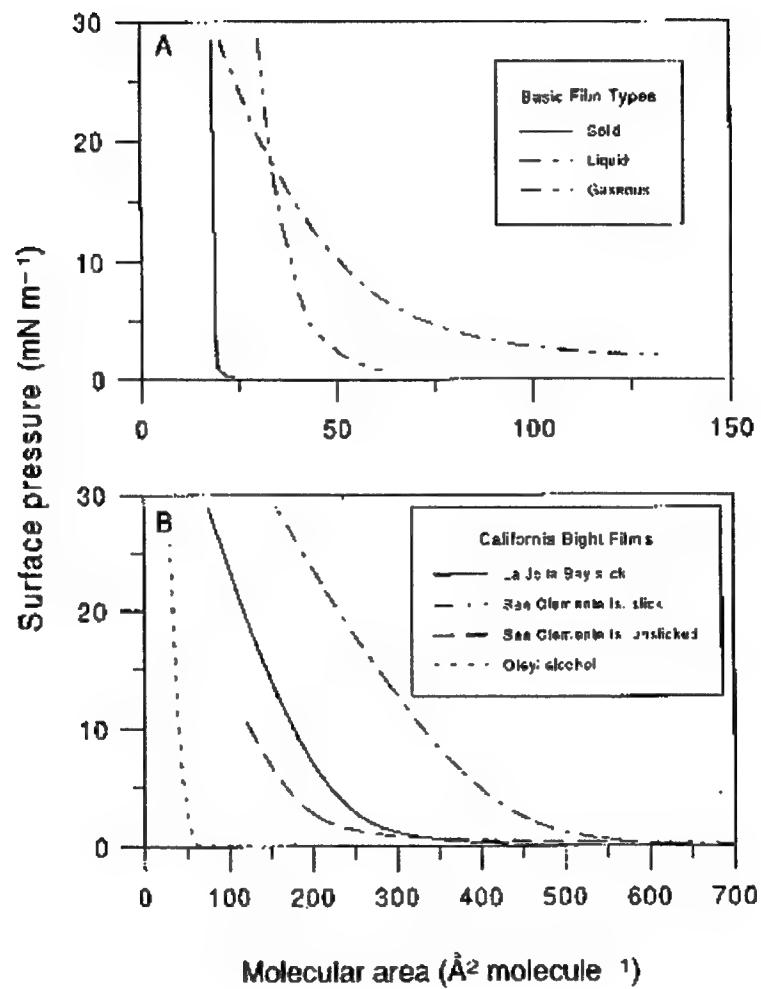


Figure 7-10: Sample π - A curves for (a) three basic monolayer film types and (b) representative sea-surface microlayer films from the California Bight. Adapted from Frew and Nelson [37]

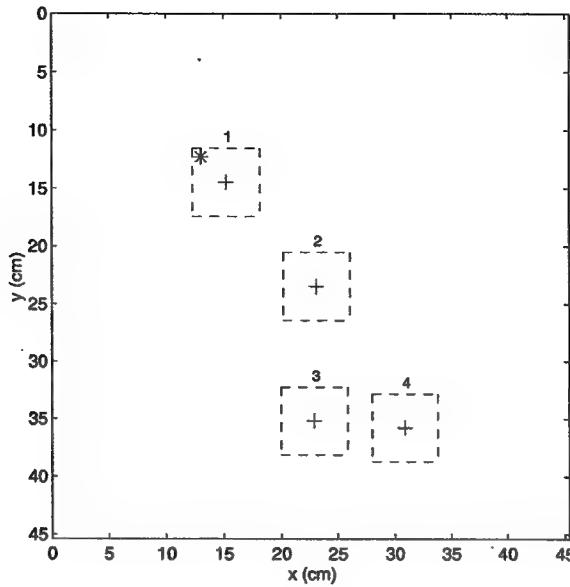


Figure 7-11: Locations of the surface PIV measurement locations. The dashed boxes indicate the $648 \text{ px} \times 648 \text{ px}$ imaged area (approximately 6 cm square). The locations are referred to in this thesis by the numbers above each region. The * marks the location of the time series PIV point measurements.

7.4.1 Experimental approach

The flow field at the free surface of the grid tank was explored using the surface PIV technique described in detail in section 4.2. As alluded to in the previous chapter, mapping the complete plan section of the surface is prohibitively time-consuming; additionally, maintaining consistent chemical surface properties over such a long time would be questionable, inducing further complications. In lieu of a complete mapping, four locations were specified for study (shown in figure 7-11). These locations are referred to in this thesis by number: 1, 2, 3, 4. Two imaging schemes were employed for the surface PIV measurements. In the first, the ensemble mode, a square region of size $648 \text{ px} \times 648 \text{ px}$ was sampled by the CCD camera. This image size corresponded to approximately 6 cm square. Two frames were collected every 5 seconds for approximately 16 minutes. This yielded 192 frame pairs $\Rightarrow 192$ velocity fields used for ensemble averaging. The processing of the ensemble pairs was held constant for nearly all the grid turbulence data collected. Hybrid correlation was employed throughout. The subimage size was $64 \text{ px} \times 64 \text{ px}$ (5.8 mm square), and

vectors were computed every 16 px (1.4 mm) in both horizontal directions. For one data set, later referred to as *dirty*, the subimage size was 32 px \times 32 px (2.9 mm square). Outlier removal was enabled, but no further post-processing was performed. In the second imaging scheme, the time series mode, the camera was configured to collect a 160 px \times 160 px region. Using this reduced image size, frames were acquired at the maximum camera frame rate of 29 Hz for a little over 10 minutes. This produced 9000 velocity fields that formed a 14.5 Hz time series. Processing these images was done with the direct spatial correlation using the 64 \times 64 subimage size. By spacing the vectors 32 px in both directions, the resulting velocity field was a 3 \times 3 matrix. This matrix was used to generate a time series of velocity, using the center node, and time series of derivative quantities, such as vorticity and divergence. In either mode, concurrent gas-transfer measurements were made, yielding contemporaneous surface hydrodynamics and interfacial transport measurements.

The choice of the 64 \times 64 subimage size was necessary due to the particle seeding density in the surface images. In the previous design of the tank, where surface contamination was likely an issue and the forcing energetics were modest, the free surface experienced little renewal or bursting. Instead, the flow was comprised of an array of meandering eddies and apparent two-dimensional jets (see figure 7-18 for an example, where the surface is damped due to the presence of a surfactant). Particle seeding at the surface was excellent and processing of the image data was straightforward. The redesigned tank, on the other hand, gave rise to much more dynamic free-surface flows. Key characteristics of the redesigned grid tank were its improved chemical cleanliness and more vigorous forcing ability. These two factors resulted in very agitated and dynamic free-surface flow regimes. Surface bursting was predominant and the gas transfer was enhanced, which were the desired goals. However, when the grid was set in motion, the favorable initial particle seeding distribution at the surface would be quickly disrupted, with nearly all the particles being swept outward toward the tank walls in a short amount of time. This was a natural consequence of the tank mean circulations identified in chapter 6. Once near the tank walls, the particles tended to conglomerate there rather than being entrained back into the bulk. Some attempts were made to render the tank side walls non-wetting near the waterline (e.g., lining this area of the walls with paraffin film). However, none of the approaches were found satisfactory. The

net result of this behavior was image data with relatively low particle seeding density. In order to compensate for this, a larger subimage size was used, allowing enough particles to be included in the subimage, and producing reliable velocity estimates. A critical assessment of the surface PIV technique for this type flow is given in chapter 8.

The spatial resolution of PIV is closely related to the subimage size. Flow scales smaller than the dimensions of the subimage are not adequately resolved, even though the spacing of the velocity vectors (based on the average distance between individual particles) may be smaller, as was the case here. Such subimage overlapping only increases the number of measurements, but does not enhance the spatial resolution. For almost all of the free-surface PIV data in this chapter, the physical size of the subimages was approximately 6 mm. Flow scales smaller than this size are strongly suppressed through processing; scales near twice the size of the subimage can be sufficiently resolved (Willert and Gharib [144]). The spatial resolution of the PIV measurement will play a role in the quantitative estimates of the flow dynamics. Because of the finite size of the subimage, the contributions of the small-scale eddies to the estimate of the RMS turbulence will be neglected, resulting in underestimation. However, this underestimation should be relatively small since the smallest scales contribute the least to the RMS turbulence estimate. Simple calculations assuming a $-5/3$ spectrum of isotropic turbulence between the Kolmogorov scale and the turbulent integral scale of the largest eddies indicate that the underestimation of the RMS turbulence due to the limited PIV resolution would be about 10–20% for the present conditions. The measurement of surface divergence, which will become important later in this chapter, is also affected by the spatial resolution of PIV. Brumley and Jirka [14] have derived theoretical results showing that the smallest scales contribute the most to the surface divergence. However, they argue that it is the cumulative divergence, or strain, which accounts for the eddy time scales, that is most significant to gas transfer. In this sense, they find that the larger eddies contribute slightly more than the smaller ones, the ratio of the contributions (largest eddies to smallest) being proportional to $Re^{1/8}$, or about a factor of two for the conditions studied here. Thus, some portion of the true cumulative divergence will be lost because of the resolution of the PIV measurement. Further assessment of this loss requires more sophisticated knowledge of the actual spectra (temporal and spatial) of the divergence.

Data set	Re_{HT} included	$\bar{\pi}$ (mN/m)
<i>cleaned I</i>	282, 469, 596, 730, 751, 898, 974	-0.01 ± 0.28
<i>cleaned II</i>	282, 469, 596, 730, 898, 974	-0.16 ± 0.28
<i>dirty</i>	282, 384, 469, 596 730, 898, 974	1.16 ± 0.27 3.51 ± 0.26
<i>oleyl alcohol I</i>	898 ($c = 0.013 \mu\text{g}/\text{cm}^2$) 898 ($c = 0.138 \mu\text{g}/\text{cm}^2$)	0.24 ± 0.27 2.87 ± 0.27
<i>oleyl alcohol II</i>	898 ($c = 0.043 \mu\text{g}/\text{cm}^2$) 898 ($c = 0.061 \mu\text{g}/\text{cm}^2$) 898 ($c = 0.091 \mu\text{g}/\text{cm}^2$)	0.75 ± 0.27 1.12 ± 0.27 1.28 ± 0.27

Table 7.2: Summary of the surface PIV/gas-transfer experimental conditions. For the oleyl alcohol data sets, c refers to the estimated film surface concentration. The values for surface pressure $\bar{\pi}$ reflect the average of the individual cases within each data set.

7.4.2 Surface conditions studied

The data presented henceforth represent an assemblage of a number of different experiments in the grid tank. Data were collected under a range of Reynolds numbers and for a variety of water surface conditions. Strictly speaking, each run at an individual Reynolds number likely represents a unique surface condition due to the sensitive, and often time-dependent, nature of the water surface. In order to explore the differences between a cleaned surface and surfactant surfaces, data (both surface PIV and gas-transfer) were obtained under both types of conditions. Table 7.2 summarizes the groups of different conditions studied. The group of experiments classified as *cleaned I* were collected for a cleaned grid tank with surface aspiration occurring before individual runs. Surface PIV measurements were collected at all four locations (in the ensemble mode) for the majority of the cases in this group; for a single set of runs, only location 3 was examined. The group of experiments classified as *cleaned II* were collected in the time series mode at the location specified in figure 7-11. The *dirty* data set corresponds to a group of runs where a surfactant (“cholesterol”) was added arbitrarily to the surface to inhibit observed surface bursting and renewal. (The details of the surfactant addition are discussed below.) Ensemble mode surface PIV was obtained at all four locations for this set. The *oleyl alcohol I* data correspond to two surfactant-adsorbed surfaces, the estimated surface concentrations denoted in the table. Similarly, the

oleyl alcohol II data reflect three surfactant-adsorbed surfaces. For both oleyl alcohol data sets, ensemble mode surface PIV was collected only at location 1.

The surfactants used in these experiments, “cholesterol” and oleyl alcohol, have already been discussed on page 225. In order to spread films of these surfactants, each sample was diluted in hexane. Using a precision micro-pipette, the desired amount of solution was carefully introduced at the water surface in the center region of the grid tank. Five minutes were allowed for the hexane to evaporate before experiments began. Once the run at this surface film concentration was finished, an additional amount of surfactant was added in the same manner to yield another (higher) surface concentration, and the next run began.

7.4.3 Surface hydrodynamics for cleaned and surfactant surfaces

The data reduction scheme used to analyze the surface PIV results was as follows. For the ensemble mode data, at each location studied a mean surface flow $(U(x, y), V(x, y))$ was determined by ensemble averaging the set of individual velocity fields $(\tilde{u}_n(x, y), \tilde{v}_n(x, y); n = 1 \dots 192)$ for a single run. This mean flow was subtracted from each velocity field to yield the fluctuating velocity fields for the run. Characteristic RMS surface velocity fluctuations for each run were defined as

$$u = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{N_x N_y} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} [\tilde{u}_n(i, j) - U(i, j)]^2 \right)^{1/2} \quad (7.2)$$

$$v = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{N_x N_y} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} [\tilde{v}_n(i, j) - V(i, j)]^2 \right)^{1/2}, \quad (7.3)$$

where $N = 192$, and N_x and N_y correspond to the size of the vector fields. The quantity in the large parentheses is the spatial RMS of the velocity fluctuations for each realization. Thus, u and v are just the overall mean spatial RMS fluctuations. In addition, the following two quantities were defined for convenience:

$$u_h = \frac{u + v}{2} \quad (7.4)$$

$$U_h = \frac{<|U|> + <|V|>}{2}, \quad (7.5)$$

where $\langle \cdot \rangle$ denotes the spatial mean of magnitude. To evaluate the vortical nature of the surface flow, the surface enstrophy was examined. Enstrophy, a measure of the rotational energy of the fluid, is defined as the square of the fluid vorticity, and in this case is written as

$$\varepsilon = \omega_z^2, \quad \omega_z = \partial \tilde{v} / \partial x - \partial \tilde{u} / \partial y. \quad (7.6)$$

The vorticity was computed over each velocity field (\tilde{u}, \tilde{v}) using the circulation, Γ :

$$\omega_{i,j} = \frac{\Gamma_{i,j}}{dA}, \quad (7.7)$$

where $dA = 4\Delta x \Delta y$ is the enclosed area of the circulation contour, with Δx and Δy as the spacings between adjacent velocity measurements. The local circulation at node (i, j) (positive clockwise) is given by

$$\begin{aligned} \Gamma_{i,j} = & (\tilde{u}_{i-1,j-1} + 2\tilde{u}_{i-1,j} + \tilde{u}_{i-1,j+1}) \frac{\Delta x}{2} + \\ & (\tilde{v}_{i-1,j+1} + 2\tilde{v}_{i,j+1} + \tilde{v}_{i+1,j+1}) \frac{\Delta y}{2} - \\ & (\tilde{u}_{i+1,j+1} + 2\tilde{u}_{i+1,j} + \tilde{u}_{i+1,j-1}) \frac{\Delta x}{2} - \\ & (\tilde{v}_{i+1,j-1} + 2\tilde{v}_{i,j-1} + \tilde{v}_{i-1,j-1}) \frac{\Delta y}{2}. \end{aligned} \quad (7.8)$$

The surface divergence, introduced in chapter 2 as being an important process in air-water transport, was also computed over each velocity field (\tilde{u}, \tilde{v}) using a first-order central difference:

$$a_{i,j} = \frac{\tilde{u}_{i+1,j} - \tilde{u}_{i-1,j}}{2\Delta x} + \frac{\tilde{v}_{i,j+1} - \tilde{v}_{i,j-1}}{2\Delta y}. \quad (7.9)$$

For the two quantities just outlined, ε and a , spatial averages over each two-dimensional realization were taken; for the divergence, it was the *magnitude* that was averaged. Each ensemble of 192 values then was averaged to yield a single characteristic value at each measurement location for each run. For the surface PIV data collected in the time series mode, the same quantities discussed above were computed. The time series means, U and V , were used to compute the fluctuating components, and from the fluctuations, the RMS values u and v were determined. For the enstrophy, a simple time-average yielded the

characteristic value for a particular run. For the divergence, the RMS was estimated by a direct calculation from the time series as well as through the integrated spectrum (see equation 2.55); both methods yielded values within a few percent of one another. The RMS estimate of the divergence was chosen in light of the ideas of McCready *et al.* [92] as discussed in section 2.4.3.

The hydrodynamics of the surface flows were investigated in a manner similar to the bulk flow in chapter 6. Measurement uncertainty for the PIV derived results is treated in appendix B and will be indicated in the subsequent figures with vertical (or horizontal) bars. To first assess the turbulence levels relative to the local mean flows at the surface, the characteristic ratio u_h/U_h was investigated (figure 7-12). In general, at locations 1, 3, and 4, the magnitude of the fluctuating velocities is roughly twice that of the mean flow velocities. At location 2, the mean velocities are a smaller fraction of the fluctuations, likely indicating that the mean flow in this region is weaker. The data in figure 7-12, compared to that in figure 6-15 where the ratios were on the order of 4, imply that the turbulence decays faster away from the grid than does the mean flow. There does not appear to be any clearly discernible dependence on the surface condition or on bulk turbulent Reynolds number.

Figure 7-13 shows the dependence of the surface turbulence on the Reynolds number. To identify the *oleyl alcohol I, II* data points, u_h is inversely correlated with π (or, equivalently, c). This relationship holds true for all of the quantities (ε, a, k , etc.) appearing in the figures in this chapter. Also included in this figure are the predictions of Hopfinger-Toly and Hunt-Graham discussed in chapter 6. These curves are shown only for reference; fair comparisons between the predictions and the data cannot be made since the former are derived for ideal, clean interfaces. The effect of the surface condition is evident in this figure. The *cleaned I, II* data sets show fair agreement with the predicted trends. However, most of the cases where a surfactant was present fall notably below. The added stress at the interface due to the surface film inhibits in-plane motions and reduces the turbulent intensity in the interfacial region. Location 2 tends to show the lowest levels of surface turbulence of the regions studied, further indicating that the center region of the tank is an area of reduced activity.

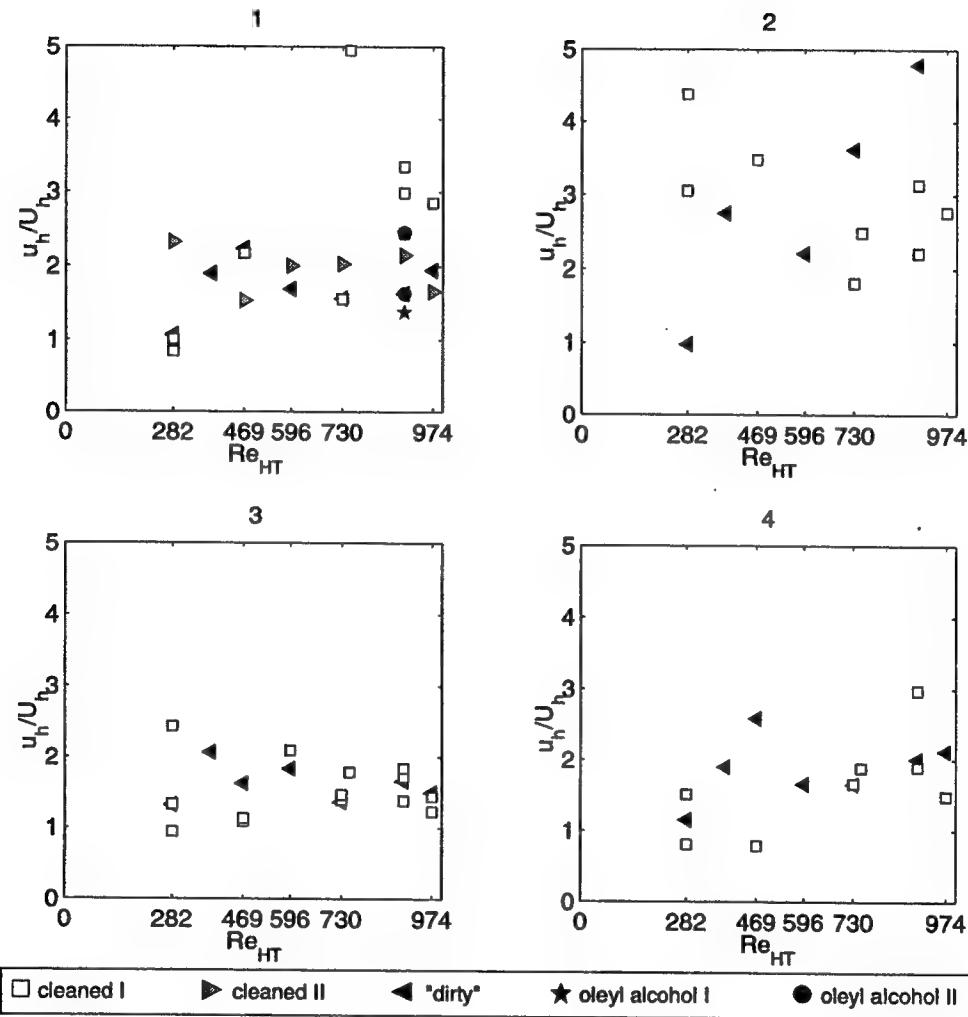


Figure 7-12: Ratios of the surface RMS fluctuating velocities to the surface mean velocities (u_h/U_h) at the four measurement locations. The locations are indicated by the numbers above each panel. The figure key specifies the respective data sets, which are discussed in the text.

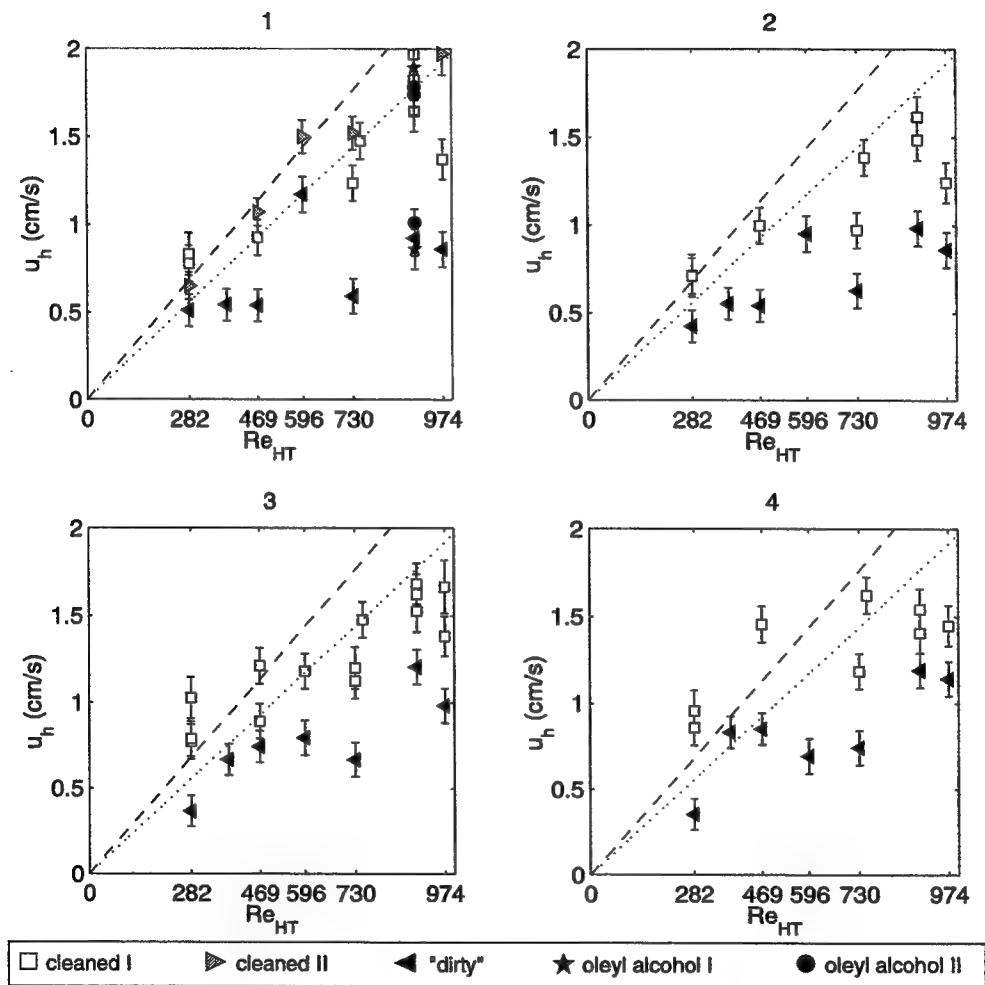


Figure 7-13: Horizontal surface RMS fluctuating velocities (u_h) at the four measurement locations. The vertical bars indicate the measurement uncertainty as discussed in appendix B. The dotted line corresponds to the Hopfinger-Toly estimate at the surface; the dashed line corresponds to the combined Hopfinger-Toly/Hunt-Graham estimate.

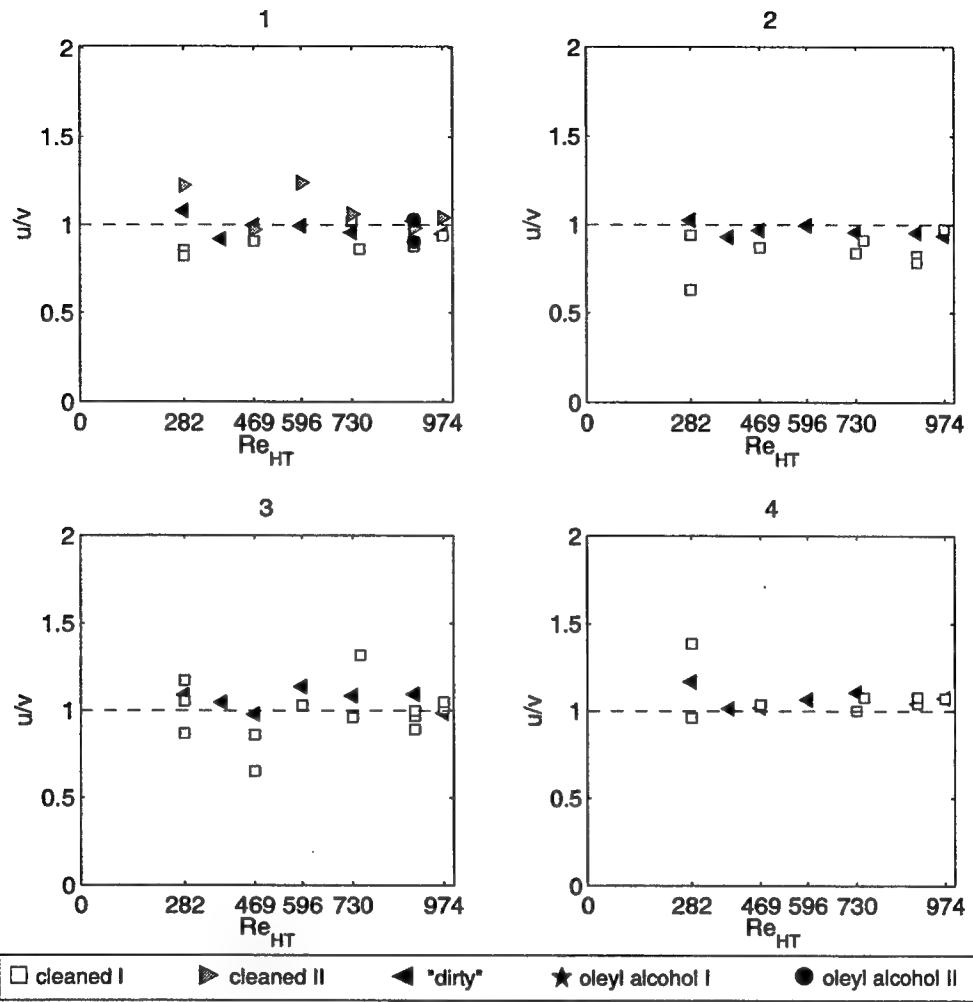


Figure 7-14: Isotropy of the surface RMS fluctuating velocities (u/v) at the four measurement locations. The dashed line indicates perfect two-dimensional isotropy.

To assess the degree of flow isotropy at the surface, the characteristic ratio u/v was computed and is shown in figure 7-14. It does not appear that the surface condition noticeably influences the isotropy. The Reynolds number also is observed to have no effect. In general, the data are scattered about the perfect isotropy measure ($u/v = 1$), with very slight systematic deviation at locations 2 and 4.

The characteristic surface enstrophy is shown in figure 7-15. As with the surface turbulence, the rotational energy shows a clear sensitivity to the presence of a surface film. Values are reduced noticeably between the *cleaned I* and *dirty* data sets, for example. The

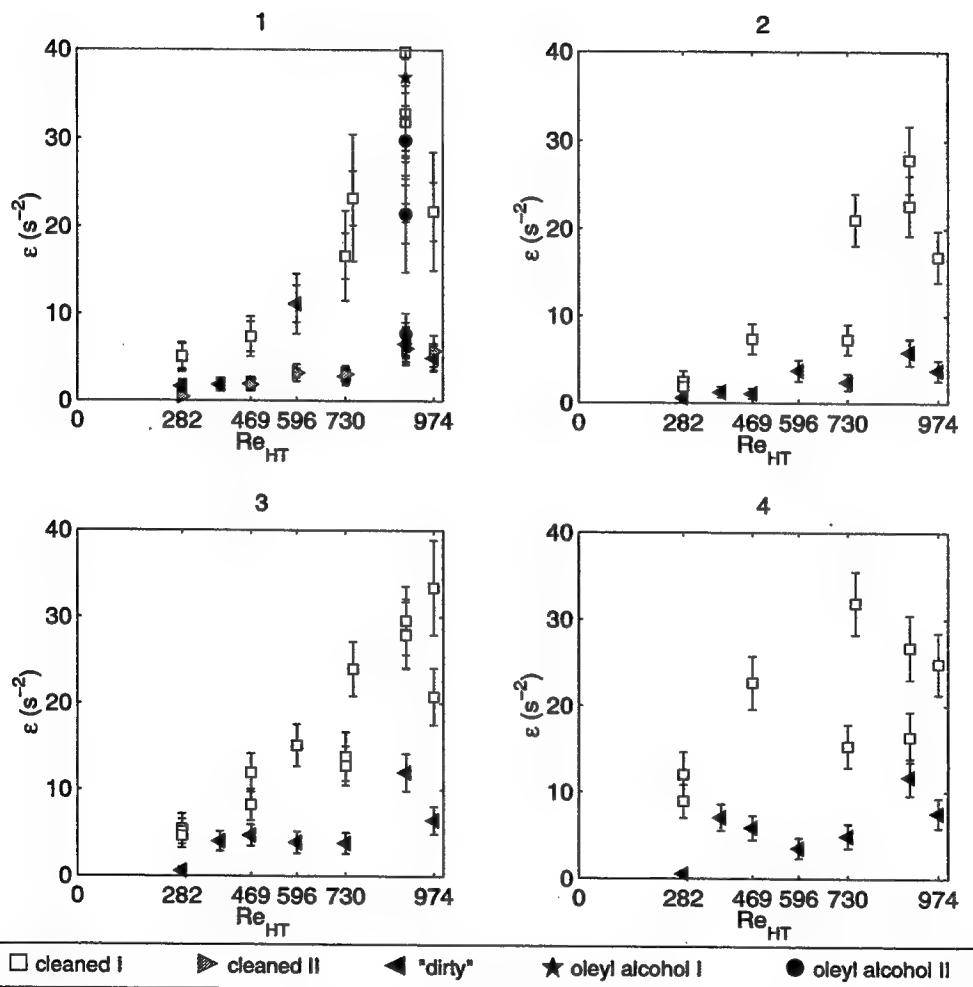


Figure 7-15: Surface enstrophy (ε) at the four measurement locations. The vertical bars indicate the measurement uncertainty as discussed in appendix B.

enstrophy levels appear relatively uniform over all locations, and generally increase with Reynolds number, although for the *dirty* data, the Reynolds number dependence is weak.

Finally, the characteristic surface divergence is shown in figure 7-16. Again, surface condition plays a significant role. The presence of the surfactant damps the surface divergence at all Reynolds numbers. The reduction at location 2 is dramatic, effectively eliminating the surface divergence at all Reynolds numbers. Although there is much scatter, the distribution of the divergence appears relatively homogeneous, and generally increases with Reynolds number.

In order to focus more closely on the impact of a surfactant on the free-surface dynamics, the data from location 1 at $Re = 898$ were examined. These data include measurements from all data sets listed in table 7.2. Figure 7-17 shows relationships between surface pressure and the turbulence/mean flow ratio, the surface RMS turbulent velocity, the surface enstrophy, and the surface divergence. The effect of increasing surface pressure on the turbulence/mean flow ratio is slight. The trend appears to indicate that the increasing surface film presence inhibits the surface turbulence somewhat more than it reduces the magnitude of the mean motions at the surface. The RMS turbulent velocity is affected noticeably for higher surface pressures. The turbulence is reduced by a factor of two for these data. A possible breakpoint in the relationship is suggested (near 1.2 mN/m, say). However, such a limited data set prevents any such speculation. Both the enstrophy and the divergence are affected by the surface film similarly. A significant reduction occurs from $0 < \pi < 1.5$ mN/m, and beyond 1.5 mN/m, the effect is relatively constant. These results point out the strong sensitivity of surface hydrodynamics to even the smallest traces of surfactant presence. Most of the effect is realized for surface pressures less than 1–2 mN/m.

7.4.4 Flow fields for cleaned and surfactant surfaces

As already seen, the impact of a surface film on free-surface hydrodynamics is significant. The entire nature of the surface flow field is modified—by the simple addition of a tangential stress in the plane of the interface. To provide a more vivid picture of this effect, two velocity field realizations at location 1 for $Re = 898$ are presented. Figure 7-18 shows a case from the *cleaned I* data set and a case from the *oleyl alcohol II* data set. The velocity field for

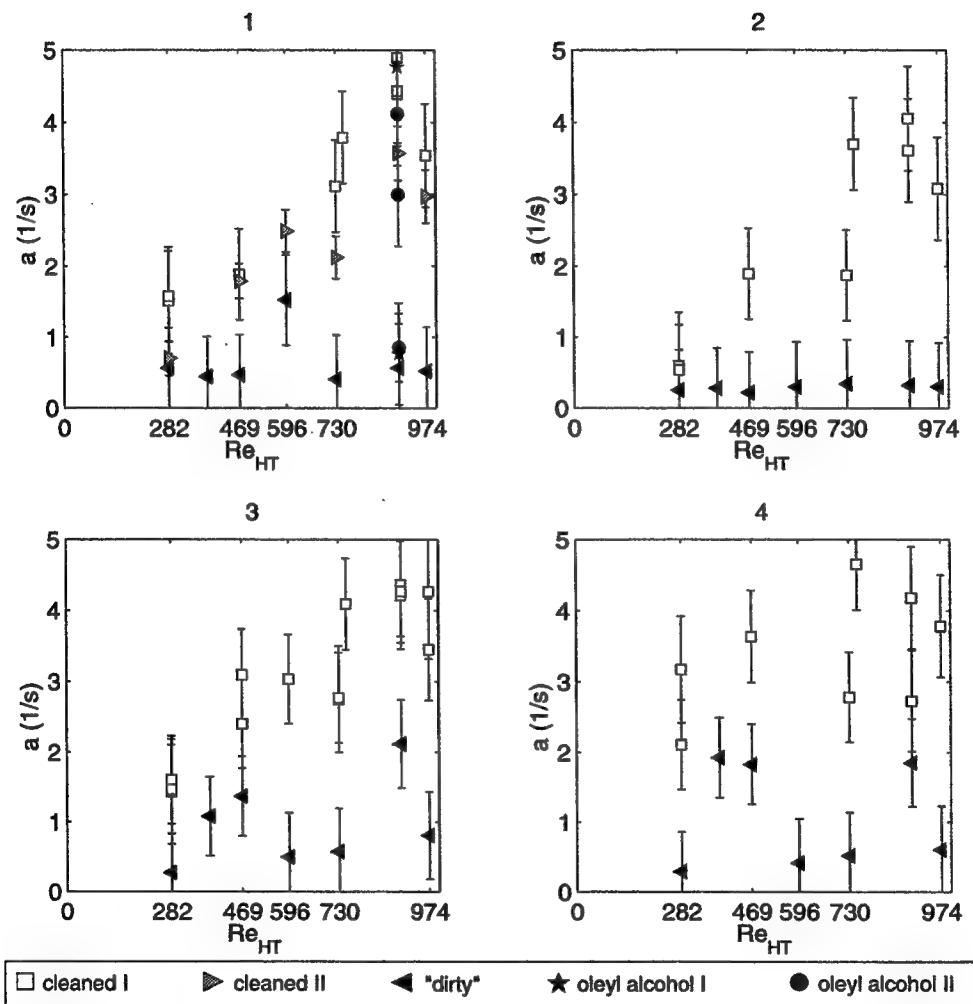


Figure 7-16: Surface divergence (a) at the four measurement locations. The vertical bars indicate the measurement uncertainty as discussed in appendix B.

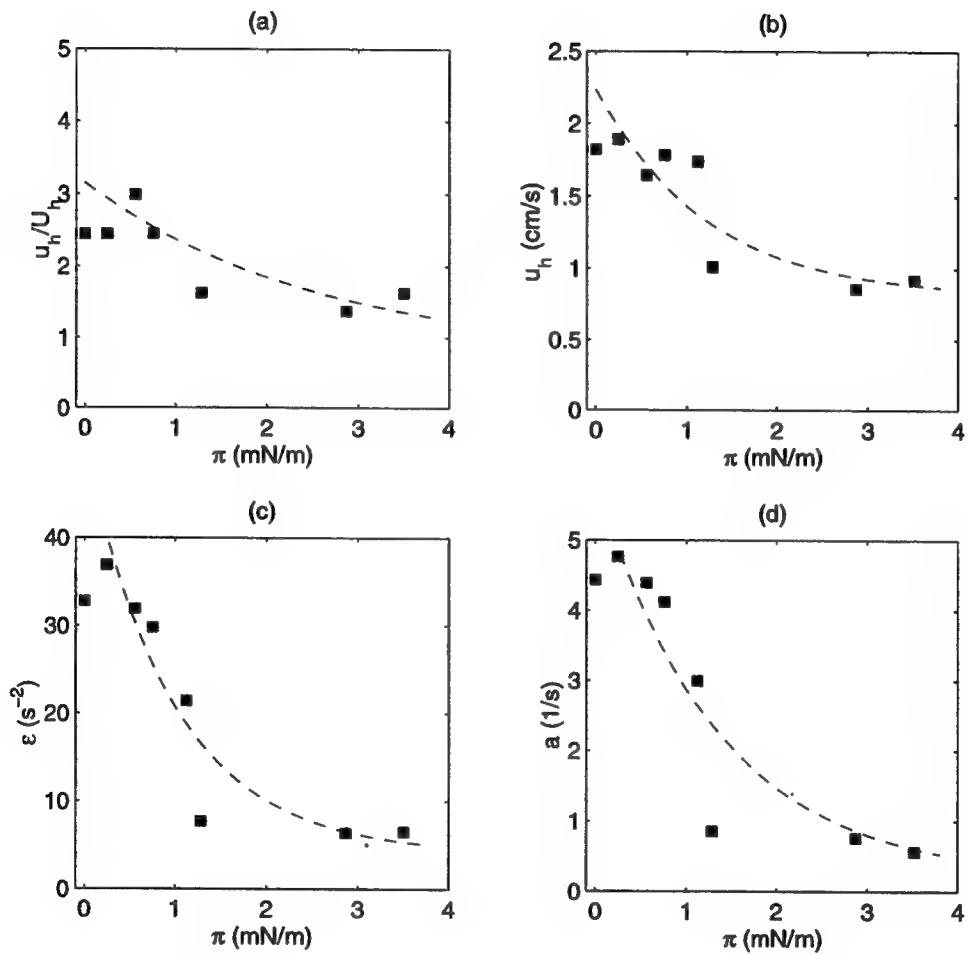


Figure 7-17: Impact of surface pressure on surface hydrodynamics. All data are for $Re = 898$ at PIV measurement location 1. (a) Ratio of turbulent velocity to mean velocity, (b) surface RMS turbulent velocity, (c) surface enstrophy, and (d) surface divergence. The dashed curves are π^{-1} trend lines.

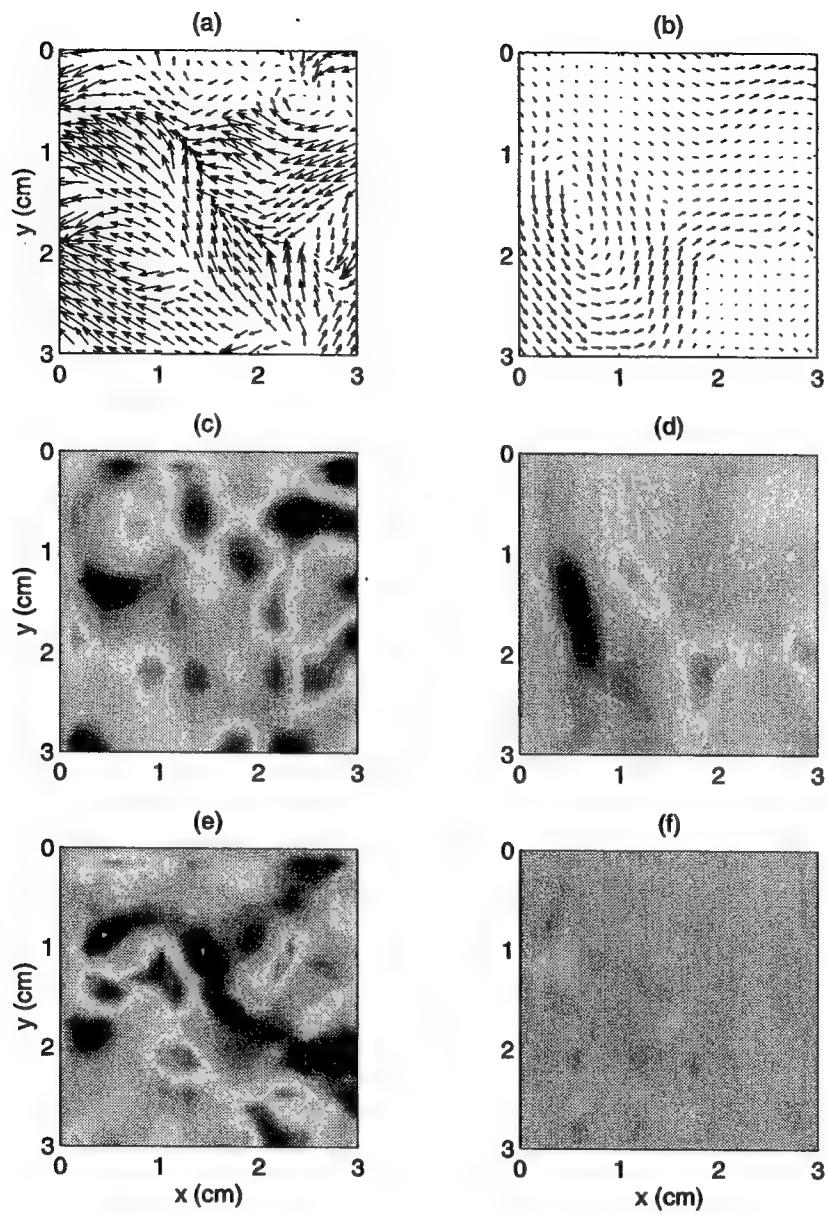


Figure 7-18: Surface flow fields for cleaned and surfactant surfaces. The left panels correspond to a *cleaned I* run at $Re = 898$. The right panels correspond to the *oleyl alcohol II* run with $c = 0.091 \mu\text{g}/\text{cm}^2$ (also at $Re = 898$). All measurements are at location 1. (a), (b) are the velocity fields; the vectors in each panel have been scaled equally. (c), (d) are contours of vorticity; the contour levels are the same for both panels: -15 s^{-1} to 15 s^{-1} in increments of 1 s^{-1} . (e), (f) are contours of divergence; the contour levels are the same for both panels: -15 s^{-1} to 20 s^{-1} in increments of 1.25 s^{-1} .

the cleaned surface condition is dominated by numerous areas of surface bursting. These motions are the result of turbulent eddies impinging at the surface in random fashion. The velocity field for the surfactant surface, on the other hand, appears less chaotic and is noticeably weaker. A large eddy motion is observed with other meandering flow elements in the vicinity. The vorticity fields are also dramatically different. In the cleaned case, the surface is populated with many local regions of both positive and negative vorticity. In the surfactant case, only two or three regions of vorticity are detectable, the most obvious being related to the large surface eddy seen in the velocity field. The difference in velocity magnitudes becomes more evident from these vorticity results. The comparison between the two surface divergence fields follows similarly. The cleaned case exhibits frequent areas of intense divergence, while the surfactant case shows very weak flow divergence. The overall sense from these two examples is that under reasonably clean conditions, the free surface is highly dynamic and energetic. The surface flow field is dominated by local regions of intense vorticity (shear) and divergence (strain). The scales associated with these shearing and straining motions are relatively small, on the order of 1 cm. When a surfactant is introduced, the surface flow becomes highly damped and the entire flow picture is modified. The total energy of the surface turbulence is reduced, in turn reducing the intensities of the vorticity and divergence. Furthermore, the flow appears to become nearly void of strong surface renewal events and is better characterized by a mosaic of meandering weak eddy motions.

7.5 Measurement of gas transfer in the grid tank

This section discusses the basis and procedures for gas-transfer measurements in the grid-stirred tank. The physical methodology is presented, the sensor is described, and aspects relevant to the measurement are discussed.

7.5.1 The transport model used for gas-transfer measurements

In this work, the gas flux through the free surface is quantified through indirect measurement of the waterside gas-transfer velocity, k . To accomplish this, the waterside mass balance for

the gas is analyzed. With the grid tank assumed well-mixed such that the gas concentration is uniform throughout the tank, the time rate of change of the gas concentration in the bulk, $C_b(t)$, is given as

$$V \frac{dC_b}{dt} = -AF, \quad (7.10)$$

where V is the water volume, A_s is the horizontal tank area, and F is the vertical flux of gas through the free surface. This flux is commonly parameterized as the product of a concentration difference and the transfer velocity (mass transfer coefficient) as

$$F = k\Delta C, \quad \Delta C = C_b - C_s, \quad (7.11)$$

written on a waterside basis. Combining equations 7.10 and 7.11, gives

$$\frac{dC_b}{dt} = -\frac{k}{H}(C_b - C_s), \quad (7.12)$$

where $A/V = 1/H$ has been used. For the case where the surface concentration is a constant, the solution to equation 7.12 is

$$C_b(t) = C_s + [C_b(0) - C_s]e^{-kt/H}. \quad (7.13)$$

This expression was used to estimate the gas-transfer velocity from time records of bulk gas concentration with appropriate steps taken to maintain a constant surface value (section 7.5.3). Specifically, an individual time series record of $C_b(t)$ is obtained (see below), and this entire record is fitted to an expression of the form given by equation 7.13 in a least-squares sense. The outcome of this fit is the estimate of the transfer velocity. Additionally, piecewise fits of the measured $C_b(t)$ were used to yield an estimate of the gas-transfer velocity as a function of time, $k = k(t)$. The length of each piecewise fit was always set at 1/3 the entire record length. The gas-transfer velocity results presented throughout this thesis are mean values of $k(t)$ over the measurement period (denoted simply as k), plus/minus one standard deviation (indicated by vertical bars in certain figures). The mean values of $k(t)$ were found to be within a few percent of the single value of k determined from the fit to the entire record. Values of k are given in units of cm/h.

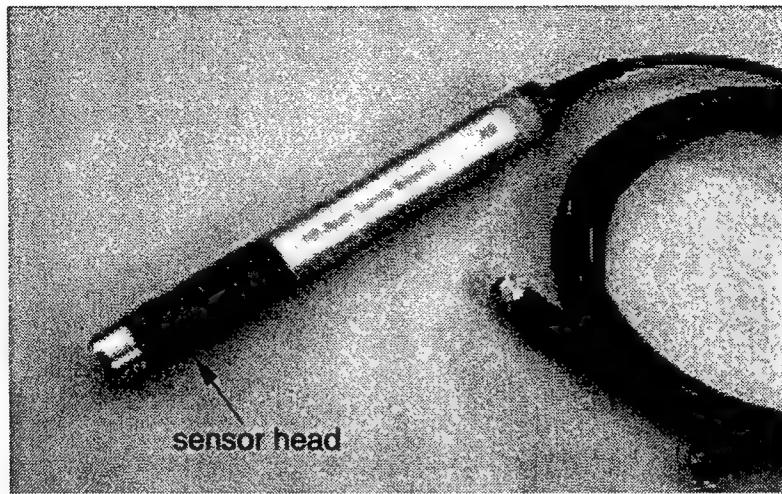


Figure 7-19: Photograph of the water quality sensor used for the bulk measurements of dissolved oxygen and temperature. The individual sensors are located in a protected head piece, indicated in the figure. The diameter of the device is 4 cm.

7.5.2 The dissolved oxygen sensor

In order to measure bulk gas concentrations, a commercially available water quality sensor was utilized (see figure 7-19). The device was outfitted with sensors for dissolved O₂, temperature, salinity, and pH. The manufacturer reports for the dissolved O₂ measurement a range of 0 to 50 mg/l, a resolution of 0.01 mg/l, and an accuracy of ± 0.2 mg/l (0 to 20 mg/l), ± 0.6 mg/l (20 to 50 mg/l). For temperature, the measurement range is -5 to 45°C, the resolution is 0.01°C, and the accuracy is 0.15°C. For the experiments in this chapter, the O₂ concentration ranged from about 5 mg/l to 25 mg/l, and the water temperature was nominally 20°C.

The O₂ sensor was installed in the grid tank through one of the floor ports. For all experiments, unless otherwise noted, the position of the sensor was $x = 11.4$ cm, $y = 11.4$ cm. The sensor was mounted in the port using two Buna-N O-ring seals such that the head protruded above the tank bottom by roughly 12 cm. The actual location of the individual sensing elements was approximately 10 cm above the tank bottom. For all data in this thesis, the sensor was used at 0.1 Hz.

7.5.3 Gas-transfer measurement procedures and diagnostics

All of the O₂ gas-transfer measurements were conducted as evasion experiments, *viz.*, $C_s < C_b$ and the flux of gas was out of the water. Experiments were always conducted in a room temperature environment. In order to hold the surface O₂ concentration known and fixed, the tank headspace (sealed with the tank lid) was flushed continuously with nitrogen gas. This allowed the assumption $C_s \approx 0$ mg/l to be made in equation 7.13. For a small number of experimental runs, UHP N₂ gas was used. For all other experiments, gas from a liquid nitrogen (LN₂) tank was used. Special handling of the LN₂ gas was necessary to bring the temperature of the gas up to near room temperature. This entailed sending the raw LN₂ gas through an aqueous equilibration system. The process is shown and described in figure 7-20. The equilibrated N₂ gas was sent into the grid tank headspace through one of the lid gas ports. A second gas port in the lid was left open for headspace flushing. In this manner, the headspace would be purged of air (and O₂). During all grid tank gas-transfer experiments, the headspace was flushed continually.

A short experiment was conducted in order to determine the N₂ flushing time required to remove all tank headspace air and O₂. This was most easily accomplished by using an available CO₂ gas analyzer to monitor the levels of CO₂ in the tank headspace during nitrogen flushing. Since the diffusivities of the two gases are similar, CO₂ gas was considered a fair proxy for O₂ gas. The CO₂ analyzer was set up to sample the outflow from the tank headspace flushing port. The tank was filled to the standard water level and the lid was sealed. The equilibration system was allowed to run for approximately 30 minutes, and then the equilibrated N₂ gas was sent through the grid tank headspace. The CO₂ levels were measured from this point on. The CO₂ concentration in the headspace was found to drop 90% within 3 minutes, and was effectively zero in 10 minutes. For all gas-transfer experiments in this work, the grid tank headspace was purged for 5 minutes before commencing a run.

The estimate of the gas-transfer velocity is obtained through a numerical fit of $C_b(t)$ using equation 7.13. The robustness of this estimate can be enhanced by increasing the concentration difference, *i.e.*, $C_b - C_s \gg 1$ mg/l. At room temperature, the equilibrium concentration of O₂ in water is roughly 8–9 mg/l, giving a concentration difference of about

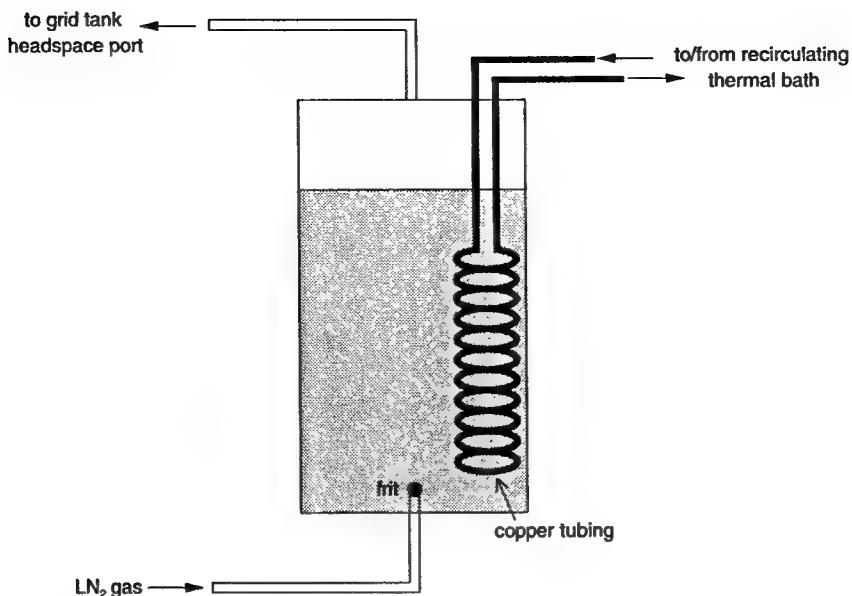


Figure 7-20: Diagram of the equilibration system used with the liquid nitrogen gas. The system consisted of a sealed cylindrical chamber filled partially with water. The water temperature in the chamber was set using a heat exchanger coil made of copper tubing. Temperature-regulated water from a thermal bath was recirculated through the copper coil. Cool LN₂ gas was bubbled into the water chamber through a porous frit located at the base. This bubbling saturated the water with nitrogen gas and purged the chamber headspace of all air leaving only N₂ gas. The N₂ gas was brought near room temperature in this manner. The warmed (and humidified) N₂ gas exited the chamber through a headspace port and was fed to the grid tank for flushing. The equilibration system was run for roughly 30+ minutes prior to grid tank experiments in order for the system to reach an equilibrium.

7 mg/l. To effect a larger surface–bulk O₂ concentration difference, the tank water was spiked with O₂ gas. This was accomplished by bubbling UHP O₂ gas through a diffusing stone temporarily placed in the tank water. The diffusing stone was made from fused crystalline alumina grains and the connecting gas line was Teflon, both items being selected so as to minimize contamination effects during this procedure. Also, the gas line and the stone always were cleaned using the 50:50 methanol/water solution and rinsed with spring water prior to insertion into the tank. The water O₂ concentration typically would be raised to about 23 mg/l with this bubbling technique. Over the course of gas-transfer experiments, the O₂ level would fall and when the concentration neared 12–15 mg/l, the bubbling procedure would be repeated.

The final issues addressed for the grid tank gas-transfer measurements were the location of the O_2 sensor and the assumption of a well-mixed system. The latter condition was initially checked qualitatively with dye visualization. A small amount of dye introduced near the water surface was found to become completely mixed for all grid forcing conditions within 10–15 seconds. An experiment, which provided additional evidence of a well-mixed system, was conducted in order to ascertain if the fixed position of the O_2 sensor in the grid tank presented any unforeseen problems. Particularly, in light of the secondary circulations in the tank, it was of interest to see if the estimated value for the transfer velocity was dependent on the location of the O_2 sensor. To test for such a dependence, measurements of the bulk oxygen concentration at different locations in the tank would be necessary. With only one O_2 sensor, this was somewhat problematic. The difficulty arose from surface film effects. If a measurement of $C_b(t)$ was made with the sensor in its standard location in the tank floor port, a measurement with the sensor in a different location would necessitate completely draining the tank so that the sensor could be removed and relocated. Upon draining and refilling, the condition of the free surface could not be considered equivalent to the previous condition and any discrepancy in the estimated value for the transfer velocity would be ambiguous—the difference could be due to location or it could be due to a slight change in surface cleanliness. With the single O_2 sensor, the following approach was adopted. An evasion experiment with the tank lid removed was conducted. This allowed the O_2 sensor to be placed in the tank from above, though the water surface, and easily repositioned during a gas-transfer run. As long as the air temperature in the vicinity of the tank remained near-constant, this approach was acceptable. The value for C_s could be estimated from published tables (e.g., Green and Carritt [45]), and while not precisely correct, could yield relative measurements of the transfer velocity at different locations, and this was all that was required.

The experiment began with the usual prefatory procedures (cleaning, filling, and aspirating). The water was spiked with UHP O_2 to a level of approximately 22.5 mg/l. The O_2 sensor was initially mounted such that the measurement of $C_b(t)$ was made at $x = 34.0$ cm, $y = 13.5$ cm, $z = 9.0$ cm (near ADV location 3). The grid was set to run at $Re = 596$; this case was chosen since it previously was found to exhibit some of the

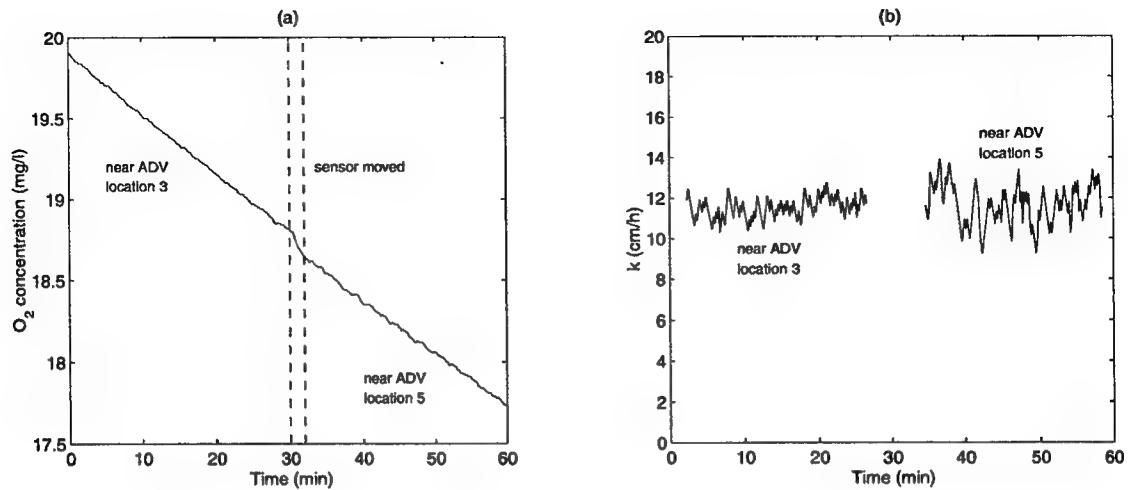


Figure 7-21: Results from the O₂ sensor relocation experiment. (a) Raw O₂ concentration data over the course of the entire experiment. The data collected for the first 30 minutes were taken near ADV location 3; the data for the last 30 minutes were taken from ADV location 5. The vertical dashed lines denote the time during which the sensor was moved. (b) Behavior of the gas-transfer velocity over time for the measurements at each location. The record length used for the sliding piecewise fits was 5 minutes. The results near $t = 30$ minutes were distorted due to the slight discontinuity in the concentration record during the sensor relocation.

most severe secondary circulations. The surface pressure measured before the start of the experiment was 0.07 ± 0.28 mN/m and the water temperature near the free surface was 17.0°C . The grid was oscillated for just over 1 hour in order to reach a steady-state transfer velocity (surface films effects were equilibrated). Data collection with the O₂ sensor then commenced and lasted for 30 minutes. The sensor then was moved to a new location, this maneuver lasting roughly 1.5 minutes. The second location for the O₂ sensor was very near the center of the tank ($x = 22.5$ cm, $y = 22.5$ cm) and at the same depth $z = 9.0$ cm (ADV location 5). Another 30 minute time series of O₂ concentration was collected at this location. At the conclusion of the run, the surface pressure was measured to be 0.19 ± 0.28 mN/m and the water temperature near the free surface was 17.5°C . The difference in O₂ equilibrium concentrations for 17.0°C and 17.5°C is only about 0.1 mg/l. The results from this experiment are presented in figure 7-21.

Several pieces of information can be obtained from these results. First, the estimates for

the transfer velocity found at the two different locations are practically identical. For the first measurement near ADV location 3, $k = 11.51 \pm 0.51$ cm/h. This estimate is taken as the mean value of $k(t)$ shown in figure 7-21b corresponding to the first measurement. The plus/minus value indicates one standard deviation. Likewise, for the second measurement near ADV location 5, $k = 11.66 \pm 0.99$ cm/h. The difference between the two mean estimates is only 1.3%. In chapter 6, ADV locations 3 and 5 were found to be quite different in terms of both mean flows and turbulence. At location 3, there is generally strong mean upwelling with peak values of turbulence intensities. Location 5, on the other hand, is generally a region of mean downwelling with relatively low turbulence levels. Because of the significant variability between these two locations and the fact that the estimated transfer velocities at each were found to be within about 1% of each other, it therefore was concluded that the location of the O_2 sensor did not impact the measurement of k . The second piece of information able to be taken from these results involves the assumption of a well-mixed system. When the sensor was moved, it appears that a concentration discontinuity resulted (see figure 7-21a). This discontinuity was estimated as approximately 0.1 mg/l, which is only 0.5%. While 0.1 mg/l is also the same difference attributable to a change in 0.5°C, it is believed that the temperature change was not the cause of the discontinuity; the discontinuity in concentration occurred over a very short period of time, while the temperature change should have occurred over the course of the two half hour long measurements. In any event, it was further concluded that the tank could be considered a well-mixed system.

A possible explanation for the observed discontinuity can be offered. To account for the slightly reduced O_2 concentration near the center of the tank, the mean flows uncovered in section 6.4.2 may be considered. For this grid forcing, one can imagine mean circulatory motions that consist of upward jets in the four corners that then recirculate back into the bulk by way of a downwelling region in the center of the tank (in addition to downwelling along all four side walls). Therefore, fluid passing by the O_2 sensor near ADV location 3 is fluid that has, in a mean sense, come up from the bulk via one of the corner jets. Similarly, fluid passing by the O_2 sensor near ADV location 5 is fluid that also has been brought up from the bulk in a corner jet, but has additionally been advected laterally near the

free surface to the central region of the tank, and is moving in a downward direction (see figure 6-21, $y = 23$ cm plane). With such a picture in mind, it can be argued that a parcel of fluid passing by the sensor at location 3 is rich in oxygen, having recently emerged from the bulk flow below. As this parcel nears the surface and is advected toward the center of the tank, oxygen gas is able to diffuse out of the parcel across the air–water interface. Upon reaching the central region of the tank, the parcel is advected downward where it eventually passes by the sensor near location 5. When the sensor at location 5 encounters this fluid parcel, it is found to be slightly depleted of O_2 gas owing to the earlier diffusion process.

7.6 Grid turbulence gas-transfer rates for cleaned conditions

To obtain a sense of the magnitudes of the transfer velocities induced by the grid tank system, experiments were performed with both grids under cleaned conditions. These experiments were exclusively gas-transfer experiments—no PIV was performed, so as to obtain upper bound estimates on the values of k for the various grid forcings. The grid tank was extensively cleaned before the group of runs, and in between each run the free surface was aspirated as described in section 6.2.4. Surface pressures (measured before and after each run) ranged from -0.18 ± 0.25 mN/m to 1.01 ± 0.24 mN/m, with a mean value of about 0.4 mN/m. The results from these experiments are included in figure 7-22. In figure 7-22a, the transfer velocity is shown as a function of the familiar Re_{HT} . In figure 7-22b, the transfer velocity is shown as a function of a new Reynolds number,

$$Re_G = \frac{SfM}{\nu}, \quad (7.14)$$

which is based solely on the grid forcing parameters. This parameter was defined in order to explore a possible relationship between an alternative measure of the grid forcing and the gas-transfer velocity. Also included in the plots of figure 7-22 are data collected from a number of gas-transfer/PIV experiments (section 7.8). Both plots show a generally increasing trend of k with Reynolds number. Interestingly, there is a fair amount of scatter in the results for these data collected under relatively similar surface conditions. It does not appear that either Reynolds number provides a superior relationship for the gas-transfer

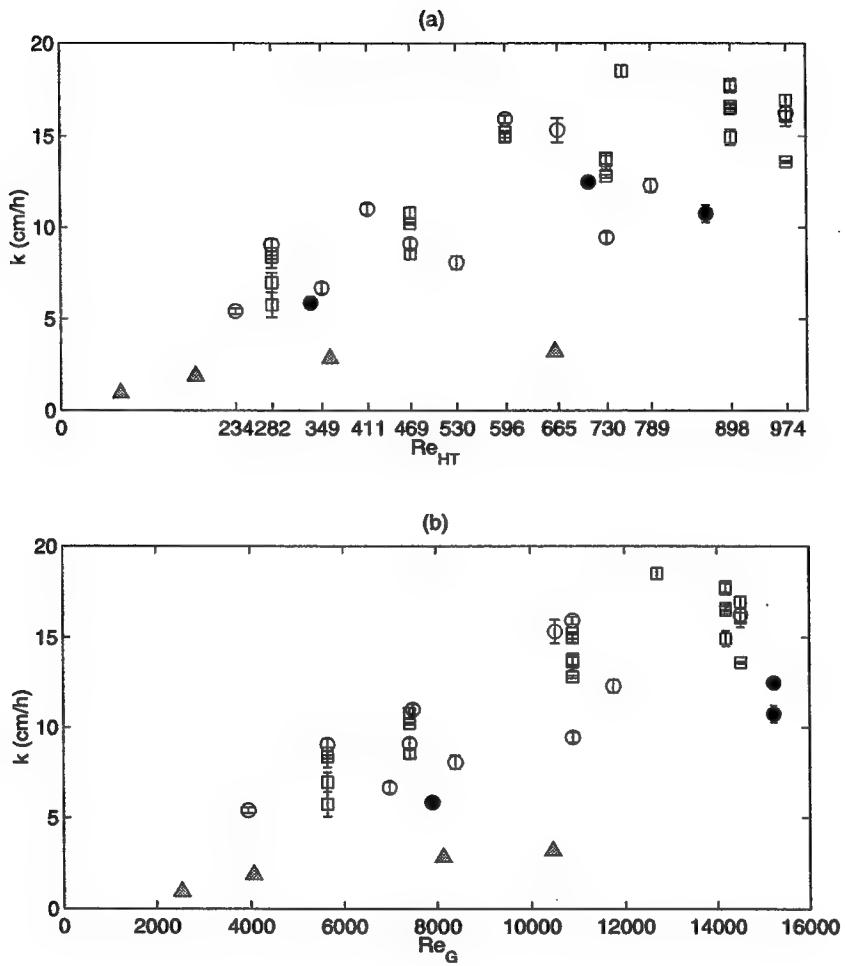


Figure 7-22: Dependence of the gas-transfer velocity on Reynolds numbers for cleaned conditions. (a) Data plotted against bulk mixing Reynolds number, Re_{HT} . (b) Data plotted against grid forcing Reynolds number, Re_G . The circles correspond to exclusive gas-transfer measurements (section 7.6); the squares correspond to contemporaneous gas-transfer and free-surface PIV measurements (section 7.8). The black markers indicate grid B, all open markers, grid A. The triangles are the measurements of Chu and Jirka [18]. The vertical bars indicate plus/minus one standard deviation about the mean $k(t)$.

velocity. Figure 7-22 additionally shows the grid-stirred tank gas-transfer measurements of Chu and Jirka [18]. These measurements were obtained in a tank 50 cm square by 40 cm deep with an 8×8 grid having the same M and d as grid A. The grid was located 25 cm from the water surface and oscillated vertically. The Chu and Jirka system therefore is relatively similar to the present system. The transfer velocity for O_2 was estimated in a manner similar to that used here, *i.e.*, by monitoring the bulk O_2 concentration over the course of *10 hours*. It is apparent that over this long period of time, the average gas-transfer velocity was considerably less than that measured for the present system. It is possible that the difference can be explained by adventitious surface films present in the Chu and Jirka experiments. The authors mention cleaning the surface, but do not report a measure of the degree of cleanliness.

The data collected during the exclusive gas-transfer experiments were used to investigate the time-dependent behavior of the gas-transfer velocity. Such behavior would be indicative of the state of surface cleanliness over time. The transfer velocity as a function of time was computed through piecewise analysis of the dissolved O_2 time series as mentioned in section 7.5.1. Sample results showing a steady $k(t)$ and a decreasing $k(t)$ are given in figure 7-23. For the steady case, the surface pressures measured before and after the run were nearly equivalent. For the decreasing case, a rise in surface pressure was detected. The latter result is likely due to surface-active material adsorbing at the interface over the course of the run. Such time-dependent behavior of $k(t)$ was often observed. Typically, the transfer velocity was measured to be greatest at the start of an individual run and at the beginning of a set of runs at the same grid forcing. For example, on one occasion, two runs at $Re = 974$ were made one after the other. For the first, $k = 16.24 \pm 0.35$ cm/h, and for the second $k = 14.66 \pm 1.13$ cm/h. The large standard deviation in the second measurement was due to a significant reduction in k over the course of the measurement. Surfactant adsorption over time is virtually inevitable. If any surface-active impurities exist when the tank is initially filled, they will make their way to the water surface over time.

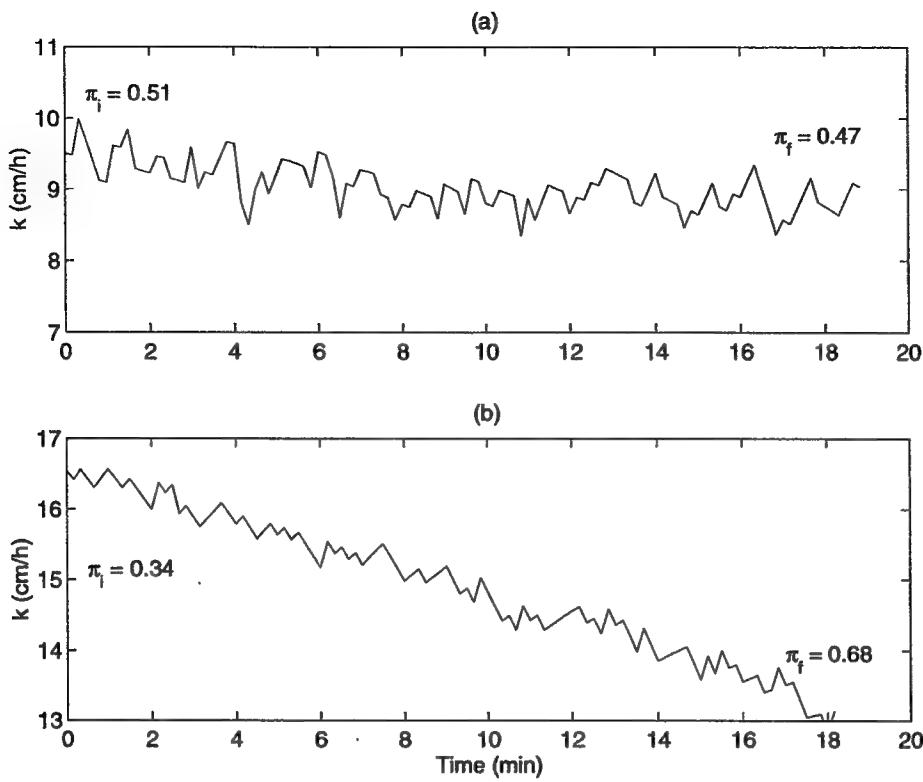


Figure 7-23: Examples of time-dependent behavior of the gas-transfer velocity. (a) A run at $Re = 282$ showing relatively steady transport. (b) A run at $Re = 974$ showing an unsteady, decreasing transfer velocity. Measurements of surface pressure (in mN/m, ± 0.25 mN/m) made before the start of each run, π_i , and after each run, π_f , are indicated.

7.7 Effect of PIV particle seeding on gas-transfer velocity

It is now appropriate to examine the impact of the presence of PIV seed particles on gas transfer. Initial experiments in the grid tank prior to redesign revealed intriguing gas-transfer results. Over the course of a number of experiments, with and without PIV measurements, it became evident that addition of a PIV seed particle mixture noticeably reduced the gas-transfer rate. The basic results are summarized in figure 7-24. In this figure, three cases are included: measurements where no PIV particles were introduced, measurements where the PIV particle mixture directly from the manufacturer was added, and measurements where PIV particles that had been rinsed were added. The rinsing involved a procedure similar to that described in section 7.3.2, but with a much simpler

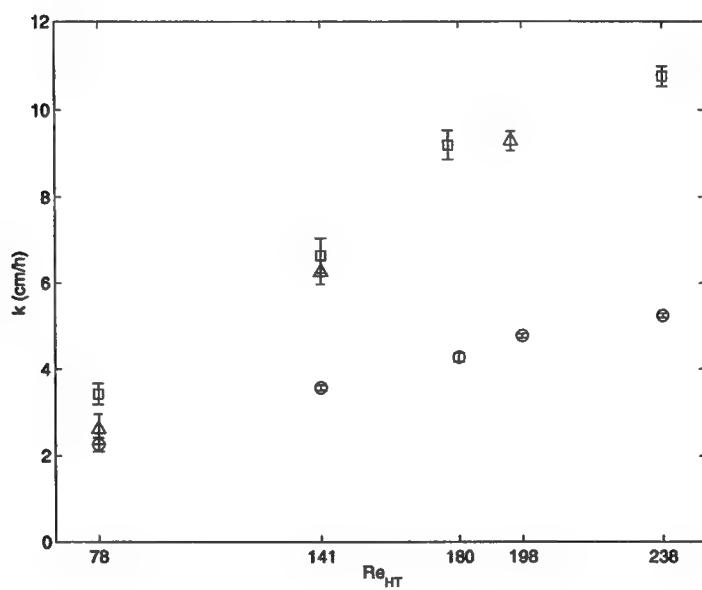


Figure 7-24: Impact of PIV seed particles on gas-transfer rates. These data are from experiments performed in the grid tank prior to redesign. Squares correspond to experiments without PIV particles, circles correspond to experiments with uncleaned, raw PIV particles added, and triangles correspond to experiments with rinsed PIV particles.

setup and using only spring water as the solvent. As the figure shows, the raw PIV particles dramatically reduced the interfacial gas exchange in the grid tank. This strongly implied that the mixture contained surface-active materials. The spring water rinsing procedure was found to be effective in cleaning the particles—the transfer velocities obtained for the rinsed particles were near those for conditions with no particle addition. These results were the impetus for properly cleaning all PIV particles used in the experiments discussed in this chapter.

A more recent experiment with $Re = 469$ showed that the particle cleaning procedure described in section 7.3.2 also proved effective. Oscillating the grid at this Reynolds number without particles yielded a gas-transfer velocity of 10.28 ± 0.20 cm/h. Continuing the same run, particles were added and the resulting transfer velocity was 9.72 ± 0.29 cm/h. Therefore, a slight reduction (5%) in the gas exchange after the introduction of the cleaned PIV seeding was found. The reduction could have been due entirely to the PIV particles, or it could have been partially due to other ambient surface-active substances adsorbing at the surface. To provide perspective for a 5% reduction, a different experiment using spring water at $Re = 469$ yielded a transfer velocity, $k = 9.1 \pm 0.30$ (figure 7-22); at the same forcing, using laboratory tap water, $k = 3.35 \pm 0.26$, a reduction of 63%. Similarly, at $Re = 974$, the spring water result was $k = 16.24 \pm 0.35$ (figure 7-22), and the result with tap water was $k = 4.76 \pm 0.25$, a reduction of 71%. In light of such results, it would appear that a potential 5% reduction due to PIV seeding is relatively insignificant.

7.8 Grid-stirred turbulence and air–water gas transfer

In this section, the surface PIV data of section 7.4 is coupled to the contemporaneous gas-transfer measurements that were made. The data sets studied in that section are carried over into this discussion.

7.8.1 Typical experimental procedure

Thus far, the experimental procedures for surface PIV, gas transfer, and surface tension have been discussed separately. The data to be reported in the present section are the result of

the integration of all three components. Below, the typical experimental procedure for a PIV/gas-transfer experiment is outlined.

1. The tank and grid were cleaned thoroughly (section 6.2.4) and the tank was filled with spring water.
2. The LN₂ equilibration system was activated and run for at least 30 minutes.
3. The tank water was spiked with UHP O₂ as described in section 7.5.3.
4. The water surface was aspirated for roughly 5 minutes.
5. PIV seed particles were added.
6. The tank lid was sealed and N₂ headspace flushing begun.
7. The initial surface tension was measured using the Wilhelmy plate technique.
8. The grid forcing parameters were set and the grid was set in motion.
9. After 5+ minutes, data collection with the dissolved O₂ sensor and the surface PIV system at the first location were begun simultaneously.
10. After a PIV data sequence was taken (\approx 16 minutes), O₂ data collection would continue until the level of O₂ had dropped more than roughly 1 mg/l since the run began. At that point, the acquired O₂ time series for that PIV location was stored. The camera and laser were repositioned to the next location, and the O₂ and PIV data acquisition recommenced. Over the course of an entire run, the grid was oscillated uninterrupted.
11. Once all desired PIV measurements were made, the grid motion was stopped and the final surface tension was measured after residual surface motions had dissipated.
12. For the next run, the grid forcing would be adjusted as necessary. If the O₂ level was depleted sufficiently, the water would be re-spiked. If a surface film was to be introduced or added, it would be done at this time. The experimental procedure would resume with step 7. Once a film was introduced to the surface, no further O₂ spiking could be done.

7.8.2 Relationships between flow dynamics and gas-transfer velocity

For all the data presented in section 7.4.3, accompanying measurements of the gas-transfer velocity were made. This provided a measurement of k for the time that each location (or single point) was studied. It was found that the surface hydrodynamics did not show a great deal of variability over the four measurement locations, and the four separate measurements of the gas-transfer velocity (one during each PIV measurement) also did not vary appreciably. The latter result implies that the surfaces studied were relatively time-invariant. This further implies that any differences observed in the surface flow measurements were not the effect of surface condition variability over time, and were actual spatial variations. In the following presentation, since the mean hydrodynamic characteristics for the four different locations were not very different, those data (the characteristic values and their associated uncertainties) are now treated as averaged over all four locations. Likewise, the four estimates of the transfer velocity were averaged. Thus, for a given run, only the averaged hydrodynamic quantities and gas-transfer velocities are considered. As evident from the vertical bars in figure 7-22, the variability in each measurement of k is small; therefore, these bars are omitted from the remaining figures.

The first relationship investigated was that between the bulk turbulence and the gas-transfer velocity. The bulk turbulence was characterized using the empirical Re_{HT} , which was shown in chapter 6 to be an accurate measure of the actual turbulent mixing Reynolds number for this grid tank. Figure 7-25 shows these results. It is quite evident that, depending on the surface condition, for a given bulk turbulence level, the transfer velocity can be widely multi-valued. The presence of surface-active material is seen to reduce the transfer velocity by a factor of four in certain cases. This is not a surprising result: interfacial air-water transport is a process that is determined by the physics very near the free surface, which is where the dominant effects of surfactants manifest themselves. Bulk turbulence away from the interface is largely unaffected by surface films and therefore cannot accurately predict the levels of gas transport. This is a key result and relates back to the first figure of this thesis on page 24. Flow measurements just below the air-water interface are still incapable of indicating the nature of the free-surface hydrodynamics that are relevant to interfacial transport. Figure 7-25 also shows the model predictions of Fortescue

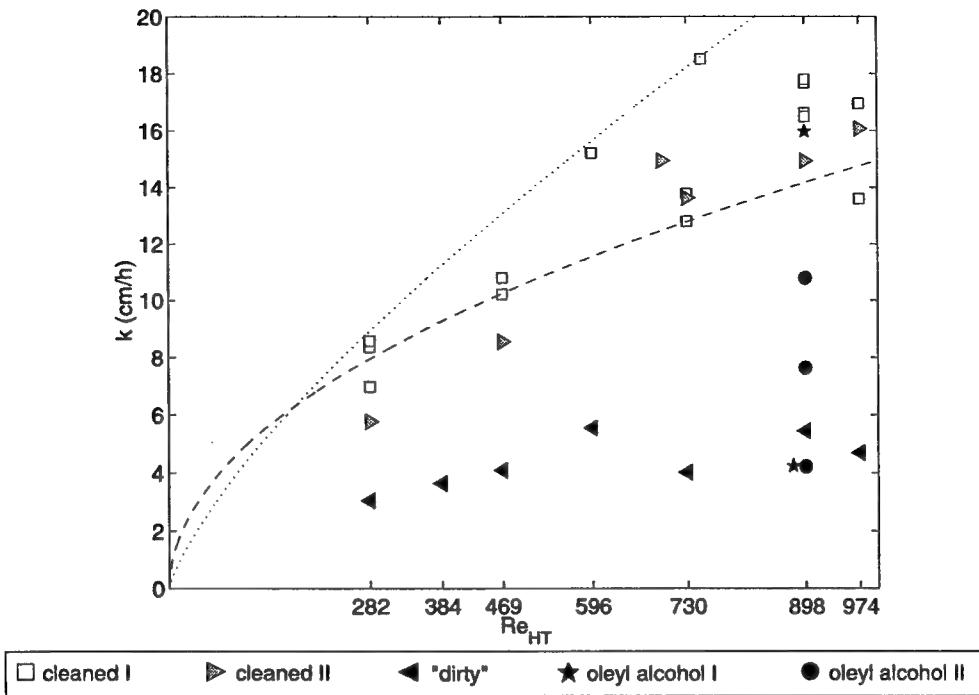


Figure 7-25: Dependence of the gas-transfer velocity on bulk mixing Reynolds number for grid-stirred turbulence. The dashed line corresponds to the large eddy model of Fortescue and Pearson (equation 2.51). The dotted line corresponds to the small eddy model of Lamont and Scott (equation 2.52).

and Pearson (equation 2.51) and Lamont and Scott (equation 2.52). In order to add these relationships to the figure of k versus Re_{HT} , it was necessary to recast the model results as:

$$k = \frac{1.46 (D\nu)^{1/2}}{2L_{HT}} Re_{HT}^{1/2} \quad (7.15)$$

for the large eddy model (F & P) and,

$$k = \frac{0.4 (D\nu)^{1/2}}{2L_{HT}} Re_{HT}^{3/4} \quad (7.16)$$

for the small eddy model (L & S). In rewriting the large eddy result, it was assumed that the RMS velocity scale $Q = u_{HT}$ and the integral length scale $L = 2L_{HT}$. To recast the small eddy result, the classical result for isotropic homogeneous turbulence was used: $\epsilon = u_o^3/\ell$; in this case it was assumed $\epsilon = u_{HT}^3/(2L_{HT})$. The two model results are very similar save for

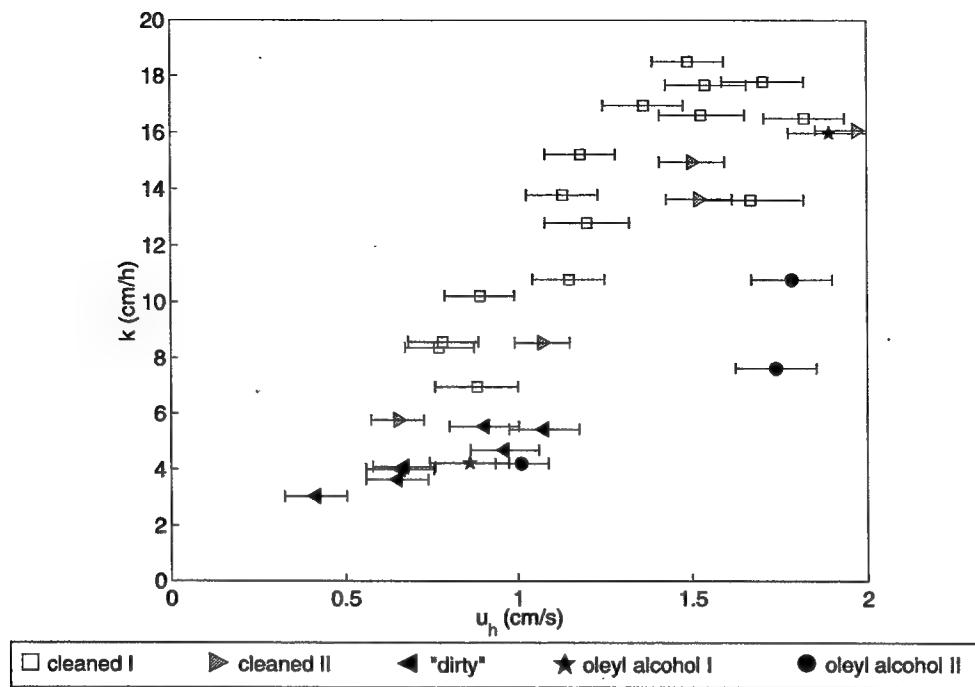


Figure 7-26: Dependence of the gas-transfer velocity on RMS surface velocity fluctuations for grid-stirred turbulence. The horizontal bars indicate the measurement uncertainty as discussed in appendix B. $R^2 = 0.62$.

the difference in Reynolds number dependence. It may be simply fortuitous, but the two curves appear to bound the observed results under cleaned conditions. However, neither model can account for the presence of a surfactant, overestimating the transfer velocity in all such cases.

The same gas-transfer data shown in figure 7-25 are now replotted in figure 7-26 as a function of the *surface* turbulence, characterized through u_h . Relating k to u_h reduces some of the scatter in the data and suggests a more clear monotonic trend. The R^2 regression statistic for these data was 0.62. However, the relationship is still somewhat degraded by surface film influences. The transfer velocities obtained for cleaned conditions are consistently higher than those under surfactant conditions. Interestingly, there may be evidence of a two-regime relationship for the measurements that include surface films (*cf.* figure 7-1). One could imagine a line drawn through the surfactant data for, say, $0 < u_h < 1.5$ cm/s and a second, much steeper line could be drawn through the three data points beyond

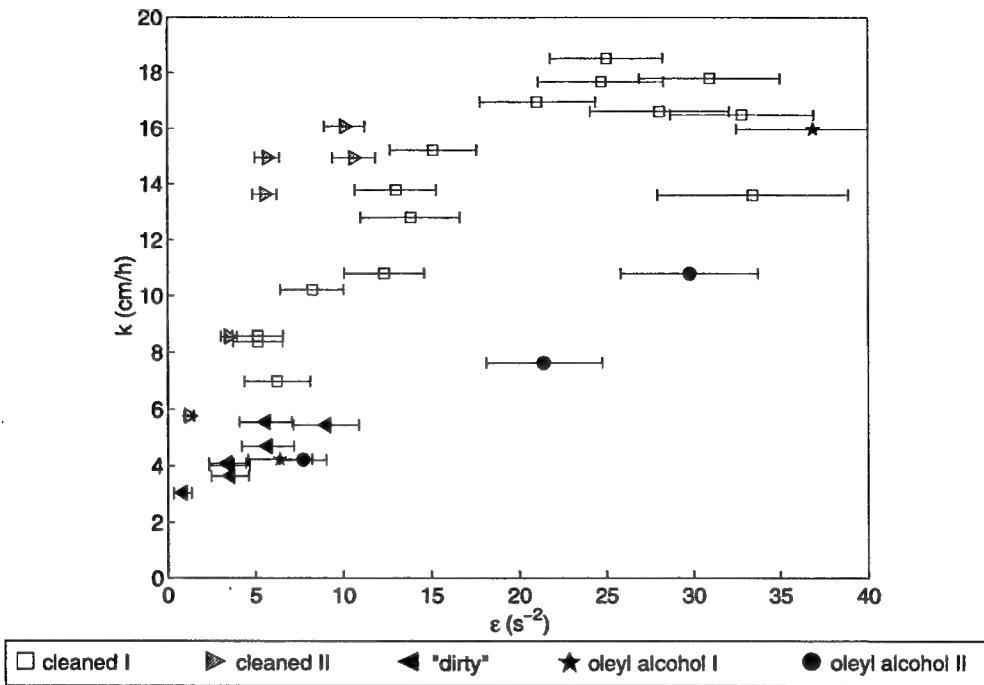


Figure 7-27: Dependence of the gas-transfer velocity on surface enstrophy for grid-stirred turbulence. The horizontal bars indicate the measurement uncertainty as discussed in appendix B. $R^2 = 0.50$.

$u_h > 1.5$ cm/s. Such a bilinear character is often observed for natural waters and contaminated laboratory systems.

Next, the same gas-transfer data are replotted against the surface enstrophy in figure 7-27. This relationship also shows a fair degree of scatter, and does not seem to improve upon a $k-u_h$ relation. The R^2 regression statistic for these data was 0.50. Again, the surfactant data show consistently lower gas-transfer velocities. Finally, to explore the role of surface divergence, the gas-transfer data are plotted as a function of $a^{1/2}$ in figure 7-28. The R^2 regression statistic for these data was 0.82, a notable improvement upon the previous relations. The model predictions of Ledwell [82] and Csanady [22] are also included in this figure. Both predictions overestimate the measured transfer velocities. This is not too surprising since both models, either implicitly or explicitly, presume *steady* divergence and renewal.

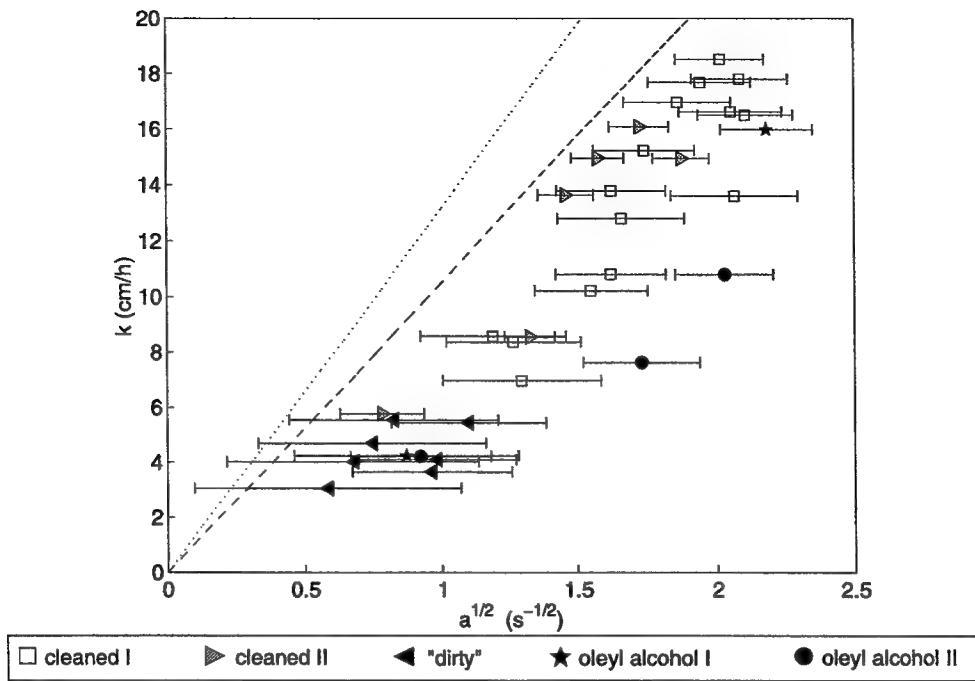


Figure 7-28: Dependence of the gas-transfer velocity on surface divergence for grid-stirred turbulence. The dashed line corresponds to the limiting result of Ledwell, $k \rightarrow (2/\pi)(Da)^{1/2}$. The dotted line corresponds to the Csanady result, $k = \sqrt{2/\pi}(Da)^{1/2}$. The horizontal bars indicate the measurement uncertainty as discussed in appendix B. $R^2 = 0.82$.

7.8.3 Grid tank surface waves and air–water gas transfer

For the case of grid-stirred turbulence, the gas-transfer velocity was found to best correlate with a measure of the free-surface divergence. This result lends confirmation to the ideas outlined in chapter 2 about the role of surface divergence on air–water gas exchange. In an attempt to further probe the robustness of such a relationship, a second type of flow was investigated: mechanically generated surface waves. A group of experiments were conducted in the grid tank (since it was optimally set up for surface PIV and gas-transfer measurements) that involved surface waves in the presence of varying concentrations of surfactant. As shown in chapter 5, the surface PIV technique developed in this thesis is capable of measuring the velocity field for undulating water surfaces with reasonable accuracy.

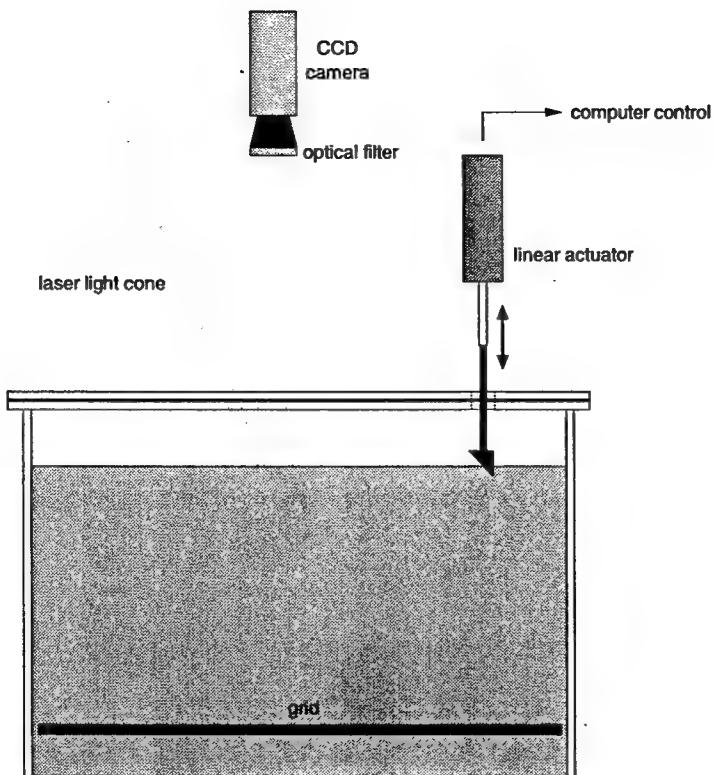


Figure 7-29: Sketch of the wave generation scheme in the grid tank. The view is from the side and only shows the upper portion of the tank. Not to scale.

Experimental arrangement and procedure

In order to easily generate waves in the grid tank, a simple plunger-type wavemaker was employed. The setup is depicted in figure 7-29. A polycarbonate wedge-shaped plunger that spanned the complete width of the tank was used to displace the free surface. The plunger was coupled to a linear actuator that was computer controlled as in chapter 5, and was programmed to follow a sinusoidal displacement with a prescribed amplitude (1.5 cm) and frequency (3.2 Hz). The plunger entered the tank lid through the hole used for the surface tension measurements. For these experiments, this hole also was used as the lone N_2 gas flush outlet. Additionally, since the fluid motions induced by the surface waves were largely restricted to the upper regions of the tank, the O_2 sensor was relocated using one of the side wall ports above the grid at an approximate position of $x = 15$ cm, $y = 23$ cm, depth = 16 cm.

In order to keep the tank water well-mixed during the grid tank wave runs, the grid was oscillated in a low energy state: $S = 10.16$ cm, $f = 0.5$ Hz. Measurement of the transfer velocity under this forcing alone without waves indicated extremely small gas-transfer enhancement ($k \approx 2$ cm/h). Surface PIV measurements were made at a single location in the center region of the tank (location 2) in the ensemble imaging mode exactly as described in section 7.4.1. The data reduction was conducted in the same way as that for the grid-stirred turbulence experiments. Five runs were performed under the same wavemaker forcing, with increasing concentrations of spread oleyl alcohol films: $c = 0.0 \mu\text{g}/\text{cm}^2$, $c = 0.032 \mu\text{g}/\text{cm}^2$, $c = 0.044 \mu\text{g}/\text{cm}^2$, $c = 0.062 \mu\text{g}/\text{cm}^2$, and $c = 0.092 \mu\text{g}/\text{cm}^2$.

Results of wave-enhanced gas transfer

The wave field that was generated using the approach shown in figure 7-29 was dominated by quasi-steady, two-dimensional standing waves with amplitudes less than 2–3 cm. While such a wave field is hardly representative of an environmental flow, it was easily generated and provided a completely different flow regime for study. To present the results from the wave experiments, figures 7-26, 7-27, and 7-28 are replotted with the new data set, identified as *waves*.

Figure 7-30 shows how the wave-induced surface velocity fluctuations relate to those obtained for the grid-stirred turbulence. Clearly, the magnitudes of the fluctuations induced by the wave field are much higher than the turbulent fluctuations created by the oscillating grid alone. Even when the surface was covered with relatively high concentrations of oleyl alcohol, the fluctuations from the waves were still higher than any observed for the grid turbulence. However, the gas-transfer velocities were not found to be correspondingly greater. In the case of the waves, the surface flow field was predominantly a sloshing wave mode in which significant surface velocities were set up. However, although appearing as velocity fluctuations, the wave orbital velocities are not turbulent and did not have the same effect on the interfacial transport. The velocity fluctuations associated with the waves studied were therefore not as effective in creating surface renewal and enhancing gas exchange. It appears that while the surface velocity fluctuations provide an improved relationship for the gas transfer when compared to a bulk turbulence estimate, such a parameterization is

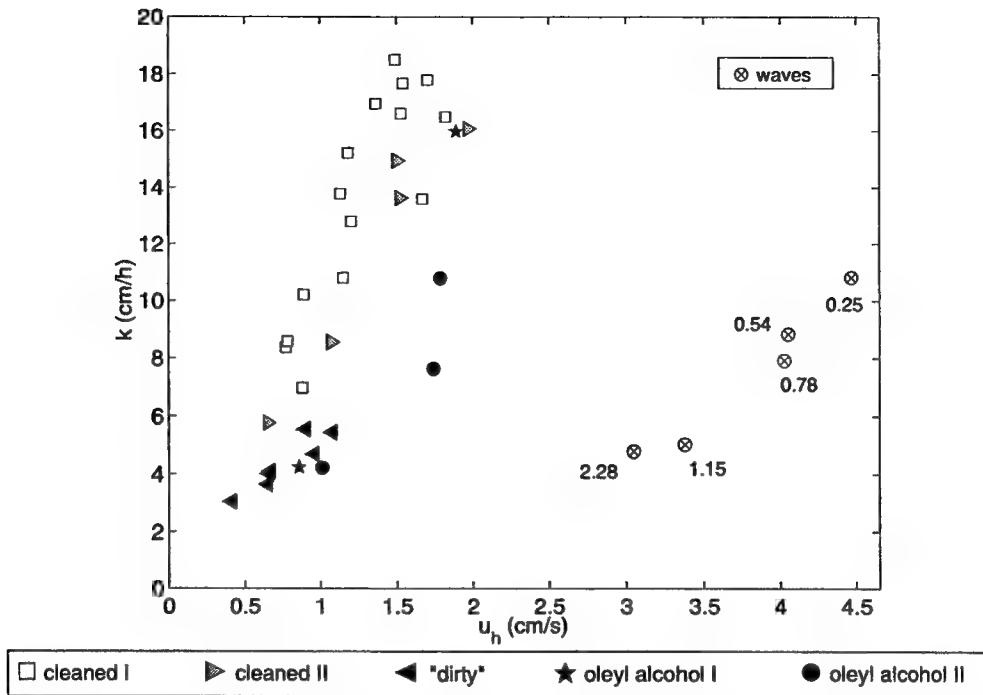


Figure 7-30: Dependence of the gas-transfer velocity on RMS surface velocity fluctuations for grid-stirred turbulence and mechanically generated surface waves. Values of surface pressure (in mN/m, ± 0.27 mN/m) for each wave run are indicated next to the markers. Increasing π corresponds directly to increasing surfactant concentration, c . The uncertainty for the grid turbulence data is as in figure 7-26; the uncertainty for the *waves* data is ± 0.31 cm/s. $R^2 = 0.02$.

dependent upon the particular flow regime. This is reflected in the R^2 regression statistic for the combined data sets of 0.02.

Surface enstrophy for both wave-induced flow and grid-stirred turbulence is shown with the corresponding measurements of k in figure 7-31. The enstrophy levels for the wave runs were extremely low, indicating that the wave field at the surface was highly irrotational as might be expected. The prevailing flow is linearly oscillatory with only weak areas of shearing motion. Inclusion of the *waves* data in the $k-\varepsilon$ relationship does little to provide any new insights. The R^2 regression statistic for the combined data sets was essentially unaffected, at 0.52.

Finally, figure 7-32 shows how the surface divergence due to waves relates to that obtained for the grid-stirred turbulence in terms of gas transfer. This result is potentially

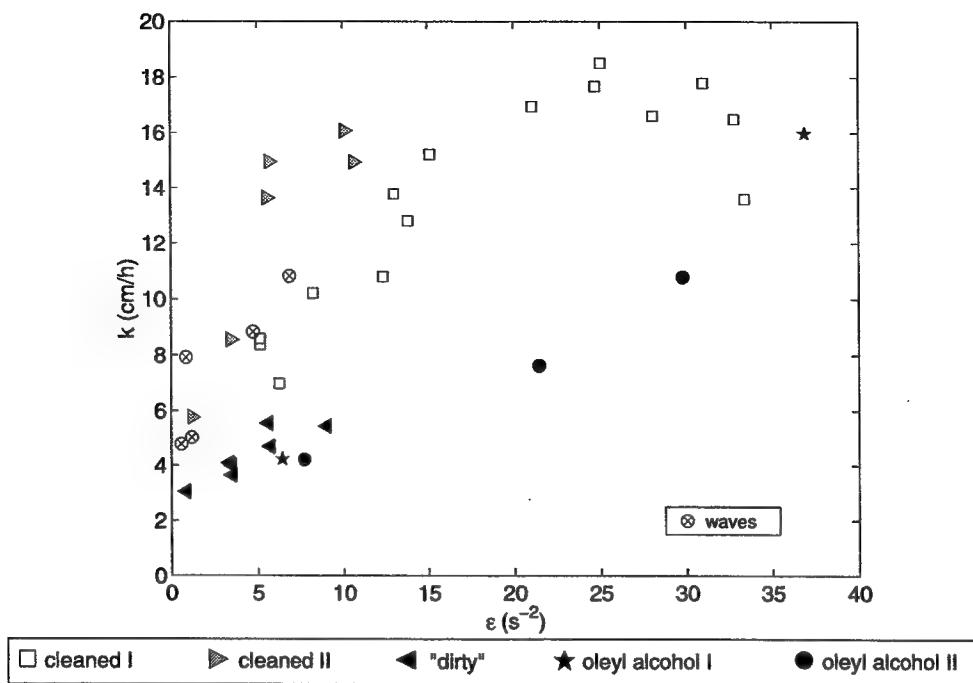


Figure 7-31: Dependence of the gas-transfer velocity on surface enstrophy for grid-stirred turbulence and mechanically generated surface waves. The uncertainty for the grid turbulence data is as in figure 7-27; the uncertainty for the *waves* data ranges from $\pm 1.4 \text{ s}^{-2}$ to $\pm 5.1 \text{ s}^{-2}$. $R^2 = 0.52$.

strong evidence supporting the concept of surface divergence as a key process in air–water gas exchange. The waves data set clusters within the data points from the grid turbulence data sets. Both flow regimes, grid turbulence and surface waves, are reconciled by the surface divergence measurement. The R^2 regression statistic for the combined data sets was 0.77—an obvious improvement over the 0.02 found from the figure 7-30 result. Thus, the surface divergence is found to effectively capture the relevant processes associated with gas transfer for two distinctly different flow regimes, and adequately account for the presence of surface films of varying concentrations. The other attempts at a relationship for k using the bulk turbulence, the surface velocity fluctuations, and the surface vorticity did not match the success achieved using the surface divergence. The shortcomings of these other measures are related to an inability to capture surface film effects and/or a lack of a direct mechanistic link to the important processes involved in air–water exchange. The

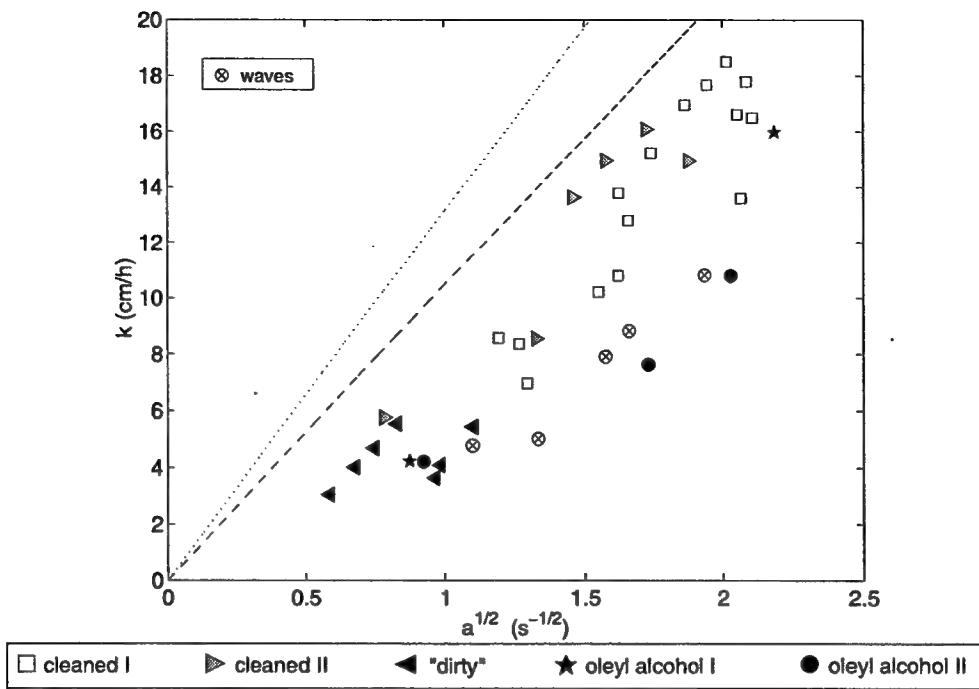


Figure 7-32: Dependence of the gas-transfer velocity on surface divergence for grid-stirred turbulence and mechanically generated surface waves. The Ledwell and Csanady curves are as before. The uncertainty for the grid turbulence data is as in figure 7-28; the uncertainty for the *waves* data ranges from $\pm 0.5 \text{ s}^{-1/2}$ to $\pm 0.9 \text{ s}^{-1/2}$. $R^2 = 0.77$.

idea that the surface divergence satisfies these two key requirements appears to be shown. This then may indicate that the surface divergence is possibly a universal process involved in air–water mass transfer.

To conclude this section, a sample flow field for a wave run is shown in figure 7-33. As in figure 7-18, the velocity, vorticity, and divergence fields are presented. The velocity field shows the organized two-dimensional wave orbital motions. Associated with these orbital velocities are vorticity and divergence. However, the magnitudes of these quantities are relatively small when compared to the grid-stirred turbulence results. Additionally, the scales of these motions are also different. For the waves, these scales are related to the wavelength of the waves in the field, which was much larger than the integral scales of the grid turbulence.

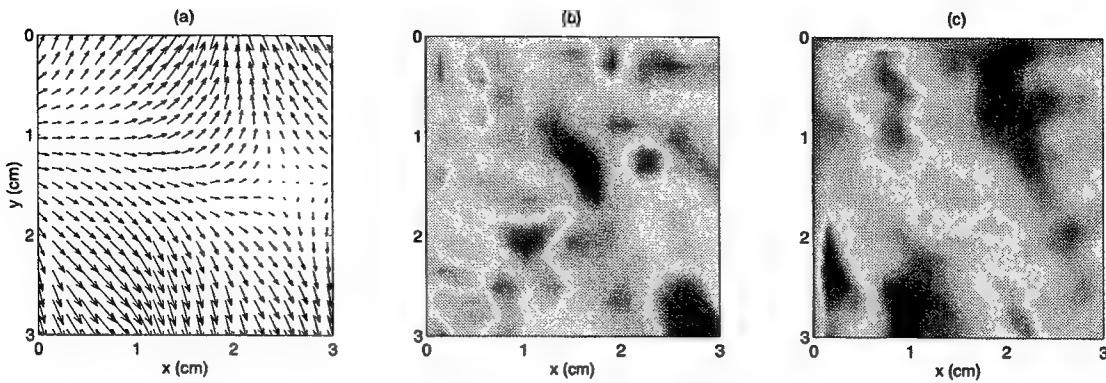


Figure 7-33: Example of a flow field from the grid tank *waves* data set. The case shown is for an oleyl alcohol concentration $c = 0.044 \mu\text{g}/\text{cm}^2$, with $\pi \approx 0.78 \text{ mN/m}$. (a) Velocity field. (b) Vorticity field; the contour levels are: -4 s^{-1} to 2 s^{-1} in increments of 0.1 s^{-1} . (c) Divergence field; the contour levels are: -4 s^{-1} to 10 s^{-1} in increments of 0.25 s^{-1} .

7.8.4 Effect of surface films on gas transfer

The *oleyl alcohol I, II* and *waves* data sets provide measurements of the gas-transfer velocity for varying levels of surface pressure for two different flow regimes. To visualize the impact of a surface film on gas transfer, figure 7-34 shows these three data sets in a k - π relationship. The fact that both the grid turbulence data for $Re_{HT} = 898$ and the wave data can appear on the same plot is purely coincidental. The trend seen in this figure is not unlike the trends found earlier for the effect of surface pressure of surface turbulence, vorticity, and divergence. The greatest impact of the surfactant is incurred at very small levels of surface pressure (less than about 1.5 mN/m); beyond 1.5 mN/m, the reduction in k remains nearly constant. Included in this figure is a result from Davies [24], who studied the effect of a surface film on damping turbulent eddies at a free surface and predicted a reduction in gas transfer. The curve shown in figure 7-34 was computed using equation 2.59 with the velocity scale chosen to be u_{HT} for $Re = 898$, the constant $b = 100$, and k_{clean} equal to the peak measured value of the group (17.8 cm/h). Changing the parameter b does not lower the curve appreciably. The qualitative trends of the measurements and the Davies results appear similar. However, a stronger surfactant effect was realized in the measurements. This, too, was found to be the case in Davies' results. Also included in the figure is an

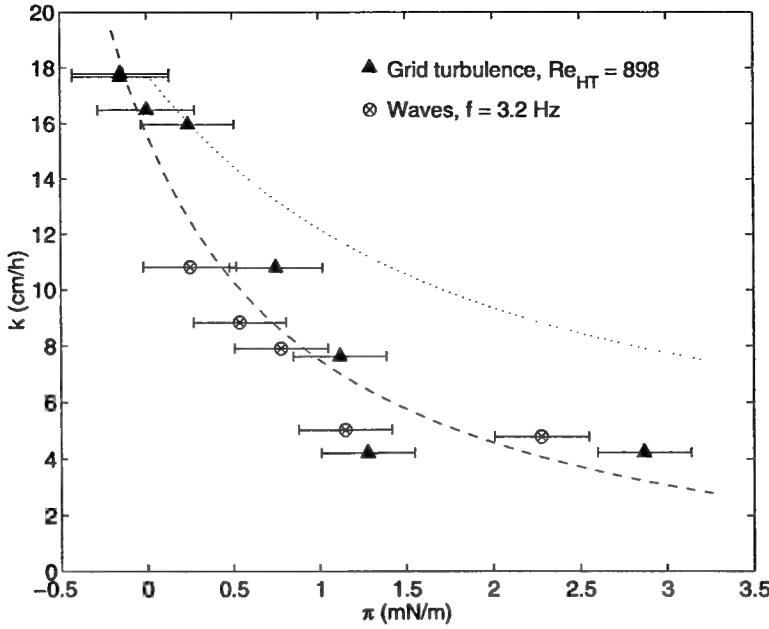


Figure 7-34: Influence of surface pressure on gas-transfer velocity for grid-stirred turbulence and mechanically generated surface waves. The frequency f corresponds to the wavemaker frequency. The horizontal bars indicate the 95% confidence intervals for the surface pressure measurements. The dashed line is an inferred π^{-1} trend line. The dotted line represents the prediction of Davies [24] as discussed in the text.

inferred π^{-1} trend line. Similar curves can be found in Frew [34] and Frew *et al.* [36] for O_2 evasion in a small stirred system. For relatively soluble surfactants, generated by a number of marine phytoplankton species, Frew and coworkers observed maximum relative reductions of the transfer velocity (*i.e.*, $k_{\text{film}}/k_{\text{clean}}$) to be 55–70%. Here, the effect is more pronounced, with $k_{\text{film}}/k_{\text{clean}}$ nearing 30%.

We can begin to understand the significance of small surface pressures on interfacial dynamics, including gas exchange, by looking at the relative contributions in the surface stress balance given by Davies in equation 2.58. In other words, how does $(2\sigma/\mathcal{R})$ compare with (π/δ_m) ? To estimate the radius of curvature, the situation for a clean surface can be considered (equation 2.57). In such case, $\mathcal{R} = 2\sigma/\rho u_o^2$, and the term $(2\sigma/\mathcal{R})$ in equation 2.58 is just of order ρu_o^2 . This is likely an upper bound on this term since when a film is present, the surface deformation should be suppressed, much like in the case of small scale waves, and the radius of curvature will be correspondingly larger. In order to evaluate the contribution

from (π/δ_m) , an estimate of δ_m is required. In section 2.2.2, a typical scaling for the viscous sublayer was noted: $\delta_m \sim LRe^{-1/2}$. Thus, $\delta_m \sim (L\nu/u_o)^{1/2}$, with L serving as a characteristic lengthscale of the turbulent eddies. The dimensionless ratio of the two contributions then becomes

$$\frac{\pi/\delta_m}{2\sigma/\mathcal{R}} = \frac{\pi}{\rho\nu^{1/2}(Lu_o^3)^{1/2}}. \quad (7.17)$$

Thus, when

$$\pi > \rho\nu^{1/2}(Lu_o^3)^{1/2} \quad (7.18)$$

the effect of a surface film can be expected to be noticeable. If we take the relevant values, $\rho = 1 \text{ g/cm}^3$, $\nu = 0.01 \text{ cm}^2/\text{s}$, $L = 2.54 \text{ cm}$, and $u_o = 1.77 \text{ cm/s}$, equation 7.18 predicts that π need only be larger than about 0.4 mN/m to begin to have an impact on the surface dynamics. This simple prediction is in excellent agreement with the observations of figure 7-34.

Further insight into the effect of small surface pressure can be found by considering the mechanism by which surface films are found to damp small scale water waves (Lucassen-Reynders and Lucassen [90]). In this situation, the enhanced damping due to the presence of a surfactant is attributed to the modification of the surface stress boundary condition. For a clean air–water interface, the simplified boundary condition for the fluid at the interface is $\mu(\partial u/\partial z) \approx 0$, and a velocity gradient at the surface is practically nonexistent. When a surfactant is present, this boundary condition becomes $\mu(\partial u/\partial z) = \partial\sigma/\partial x$, and a finite velocity gradient will exist whenever a gradient in surface tension arises. This velocity gradient is acted upon by viscosity and the energy of the wave motion is lost to the viscous dissipation $\Phi = \mu(\partial u/\partial z)^2$, resulting in wave damping. To gauge the effect of a surface tension gradient on surface flow, one could assume a variation in surface tension of about 0.5 mN/m over 1 cm, or $\partial\sigma/\partial x \approx 0.5 \text{ g/cm}\cdot\text{s}^2$. With $\mu = 0.01 \text{ g/cm}\cdot\text{s}$, the velocity gradient from the stress balance is $\partial u/\partial z \approx 50 \text{ s}^{-1}$. The dissipation (loss of energy per unit time) due to this velocity gradient is $\mu(\partial u/\partial z)^2 = 0.01(50)^2 = 25 \text{ g/cm}\cdot\text{s}^3$. A surface eddy typical of the grid-stirred turbulence studied in this thesis might have a velocity u_o on the order of 2 cm/s; the energy associated with such an eddy is $\rho u_o^2/2 = 2 \text{ g/cm}\cdot\text{s}^2$, using $\rho = 1 \text{ g/cm}^3$. This scenario implies that, on very short time scales (less than a second), the amount of

energy associated with the typical eddies would be completely dissipated by viscosity. Thus, the velocity gradient induced by the small variation in the surface tension results in highly dissipative flow, damping fluid motions near the surface.

It is useful to summarize the key aspects of our current understanding of free-surface turbulence and air–water gas exchange in terms of conceptual models, simulations, theory, and observations. Conceptualizations of the effects of surface films on free-surface turbulence have been suggested by Davies [24], Davies and Rideal [25] and Hunt [60]. Laminar interactions between a surfactant surface and a subsurface vortical flow have been studied numerically by Tsai and Yue [133], and turbulent interactions near clean and surfactant surfaces have been explored similarly in terms of scalar transport by Tsai [132]. From these studies, an understanding of the turbulence near clean and surfactant air–water interfaces emerges that relates to the experimental observations in the present study. First, the clean case may be considered. In this case, assuming a flat surface, there are two main effects of the free surface on the turbulence below: (1) an irrotational “blocking” effect that results from the kinematic boundary condition at the free surface constraining the vertical motion at the surface, and (2) a viscous effect resulting from the dynamic boundary condition at the surface that prescribes near zero tangential stress at the interface. The blocking effect is connected to the idea that eddies are restricted from approaching the surface closer than a distance equal to their size. Therefore, the effect is related directly to the eddy lengthscales and should extend over a depth roughly equal to the integral lengthscale of the turbulence (ℓ). For a clean interface, within this region, the consequence is a reduction in the turbulent vertical velocity fluctuations resulting in an increase in the horizontal velocity fluctuations through a redistribution of the turbulent energy. Shen *et al.* [119] have shown that this region also defines the depth over which the turbulent diffusivity is affected by the free surface. These effects arise from the flow adjusting to the free-surface condition that states $w = 0$ on $z = 0$, where $z = 0$ indicates the location of the free surface in a frame of reference moving with the interface. Since there are no restrictions on the horizontal components of velocity at the free surface, these turbulent components do not vary appreciably through the blockage layer. The blockage effect has been described theoretically with the source theory of Hunt and Graham [61], which treats the turbulent field as a superposition of homogeneous

turbulence and an irrotational velocity field driven by a source distribution at the interface that causes the vertical velocity fluctuations to vanish there. Numerical simulations confirm this behavior (*e.g.*, Shen *et al.* [120]; Tsai [132]), as do experimental observations like those of Brumley and Jirka [13] and the present work. For example, figure 6-23 shows the slight amplification of the horizontal velocity fluctuations at the expense of the vertical, the latter beginning to fall off at a depth near $\ell = 2L_{HT}$ as predicted. The viscous effect is realized much closer to the interface. This effect is an adjustment by the flow to the condition of zero tangential stress at the free surface, *i.e.*, $\partial u/\partial z = \partial v/\partial z = 0$ on $z = 0$. These conditions imply strong suppression of horizontal vorticity, while the vertical vorticity is not modified significantly. Therefore, the free-surface vortical motions at a clean interface are predominated by normally connected vortices (*e.g.*, figure 7-18c). The stress-free dynamic boundary condition results in a viscous sublayer of thickness $\mathcal{O}(\ell Re_o^{-1/2})$, where Re_o is based on the turbulent scales, u_o and ℓ (Brumley and Jirka [14]; Hunt [60]; Shen *et al.* [120]). Over this viscous sublayer thickness, there is typically a small change (reduction) in the horizontal velocity of an approaching eddy of $\mathcal{O}(u_o Re_o^{-1/2})$ (Hunt [60]). A similar modification occurs for the attenuation of the eddy vertical velocity, which is written as being proportional to the depth, $w \approx -(\partial u/\partial x + \partial v/\partial y)z$, in the manner shown in section 2.4.2; the increase in the attenuation rate, relative to the inviscid case, is of $\mathcal{O}(Re_o^{-1/2})$ (Hunt [60]). For $Re_o \gg 1$, these reductions are extremely small, and the viscous sublayer is of second order and negligible (Brumley and Jirka [14]). Thus, a clean free surface damps the vertical turbulent motions while effecting minimal change to the horizontal turbulent motions. The latter allow for surface divergences and vertical fluctuations in the immediate vicinity of the interface, *i.e.*, within the diffusive sublayer, which is $\mathcal{O}(\ell Re_o^{-1/2} Sc^{-1/2})$ thick. For a scalar flux out of the water, upward fluctuations bring enriched fluid toward the surface, downward fluctuations bring depleted fluid back to the bulk. This renewal of the diffusive sublayer enhances the flux at the edges of this layer through the turbulent eddy flux $w'c'$ (Coantic [19]). Well within this layer, where $w'c'$ gives way to diffusive flux, renewal events keep the concentration high close to the surface and the vertical gradient large, resulting in flux enhancement. The thickness of the diffusive sublayer, which is important for the gas flux, can be related to the thickness of the viscous sublayer through Schmidt number

scaling, $\delta_s = \delta_m Sc^{-1/2}$, where δ_s is the scalar/gas sublayer thickness and δ_m is the momentum/viscous sublayer thickness (Davies [25]). From above, the diffusive sublayer is then $\delta_s \sim \ell Sc^{-1/2} Re_o^{-1/2}$ (cf. equation 2.25).

For a surfactant surface, it is useful to consider the limiting case of a surface film that renders the interface completely immobile. The kinematic boundary condition at the interface for this situation is like that at a rigid wall where $u = v = w = 0$ on $z = 0$. As with the clean case, the $w = 0$ condition results in a blockage layer, while the condition $u = v = 0$ is manifest within a thinner viscous sublayer. The divergence at the interface is also zero, implying a reduction in the vertical fluctuations in the immediate vicinity of the interface. Instead of being proportional to z , the vertical velocity fluctuations are proportional to z^2 (viz., $w \sim z^2 u_o / (\ell^2 Re_o^{-1/2})$ (Hunt [60])), thus attenuating more rapidly when nearing a surfactant surface than when nearing a clean surface. The dynamic boundary condition at the interface is also modified, from a stress-free surface to a surface that can support a tangential stress. Again, this boundary condition influences the flow very near the surface within a viscous sublayer of thickness $\mathcal{O}(\ell Re_o^{-1/2})$ as in the clean case (Hunt [60]). However, it would be more appropriate to use a velocity scale other than u_o in calculating $Re_o^{-1/2}$, since a bulk velocity scale (say, at a depth greater than ℓ) may be largely unaffected by surfactant effects (figure 1-1 shows this clearly). Because a significant difference between the clean and immobile cases is the reduction of the horizontal velocity fluctuations across the viscous sublayer, using a surface velocity scale, for example, would be more relevant. If the surface becomes fully immobilized, the horizontal velocity near the surface can be expressed as $u \sim z u_o / (\ell Re_o^{-1/2})$ (Hunt [60]), and the reduction is of $\mathcal{O}(u_o)$, which is $Re_o^{1/2}$ times the reduction in the clean case. For interfaces that are still mobile, the present measurements indicate a maximum reduction in the surface velocity of about a factor of two (e.g., figure 7-17); the simulation results of Tsai [132] indicate peak reductions between a factor of about two to six. Based on these values, the viscous sublayer, and the diffusive sublayer, then increase to thicknesses roughly 1.5–2.5 times the thicknesses found for a clean surface. A thicker diffusive sublayer implies a reduction in k via the thin-film model ($k = D/\delta_s$). The above increases in the thickness of δ_s correlate to $k_{\text{film}}/k_{\text{clean}}$ of 40–67% based on the thin-film model; the actual measured ratios of $k_{\text{film}}/k_{\text{clean}}$ were close to 30%.

When the surface becomes immobile, associated with the dramatic reduction in the turbulent eddy horizontal velocities is a shear stress of $\mathcal{O}(\rho u_o^2 Re_o^{-1/2})$. This gives rise to strong horizontal vorticity production within the viscous sublayer; such has been shown by the simulations of Tsai [132], who also finds a reduction in the vertical enstrophy of about a factor of two for a surfactant surface. The present measurements indicate a stronger effect on the vertical enstrophy, the reduction being closer to a factor of three. This effect is, of course, dependent upon the rheological properties of the particular surfactant. Thus, the near-surface region, characterized by strong vertical gradients in the horizontal velocities, is subject to viscous dissipation, which can be quite intense, as shown previously. This highly dissipative sublayer flow damps the turbulent eddy energetics, acting like a sink layer. This can be seen as attenuating the local hydrodynamic contribution to the gas-transfer velocity parameterization through a reduction in the *local* Re_o . More importantly, the immobilized surface switches the Schmidt number dependence from 1/2 to 2/3, which will have an important effect for $Sc \gg 1$. For O_2 , the transition means a reduction in k by a factor of almost three, which is in near-quantitative agreement with the reductions observed in the present measurements. The present results indicate a factor of 3–3.5, the difference being possibly made up by the reduced local Re_o effect.

The question then becomes, to what degree does a realistic surfactant surface, which will permit some lateral interfacial mobility, resemble a rigid wall? Experiments and simulations have shown that the free-surface horizontal velocities, while not vanishing, can be reduced to 1/2 to 1/6 their clean surface values. Thus, although the rigid wall condition ($u = v = 0$) is not strictly met, the turbulence is significantly reduced at the surface of a surfactant interface and the resemblance to a rigid wall is formed. Perhaps more critical is the stress boundary condition at the surface. This condition, in its most fundamental form, expresses the balance between the viscous stress and the Marangoni stress:

$$\mu \frac{\partial u}{\partial z} = \frac{\partial \sigma}{\partial x}. \quad (7.19)$$

Therefore, the ratio of the actual Marangoni stress to a “virtual” viscous stress due to a rigid wall may provide insight into the degree to which the surface behaves like an immobile

boundary in the sense of the dynamic boundary condition. This ratio will be defined as

$$\mathcal{J} = \frac{\frac{\partial \sigma}{\partial x}}{\mu \frac{\partial u}{\partial z} \Big|_{\text{wall}}}. \quad (7.20)$$

When $\mathcal{J} = 0$, the surface behaves as a clean interface; when $\mathcal{J} = 1$, the surface resembles a rigid wall. It was noted above that the viscous stress is of $\mathcal{O}(\rho u_o^2 Re_o^{-1/2})$. A possible first approximation to the Marangoni stress is $\partial \sigma / \partial x \sim \pi / \ell$, giving the ratio

$$\mathcal{J} = \frac{\pi Re_o^{1/2}}{\rho u_o^2 \ell}, \quad (7.21)$$

which can be written in terms of a Weber number and a Reynolds number,

$$\mathcal{J} = \frac{Re_o^{1/2}}{We_\pi}, \quad (7.22)$$

where the Weber number is defined as $We_\pi = \rho u_o^2 \ell / \pi$. The choice of ℓ as the lengthscale over which the surface tension varies by π may be a conservative estimate; the surface tension may actually vary by such an amount over much smaller scales making $\partial \sigma / \partial x$ larger. Using π and ℓ in the present scaling is in line with the model of Davies as depicted in figure 2-4. For the conditions explored in this chapter, equation 7.22 indicates that π need only be 0.1 mN/m ($Re_o = 282$) to 0.6 mN/m ($Re_o = 974$) to make $\mathcal{J} = 1$. This indicates that very little $\partial \sigma / \partial x$ is required to generate an interfacial stress that is comparable to the stress at a rigid wall and that very little surface tension variation can quickly resemble a rigid wall situation, at least in an instantaneous sense. This finding is consistent with the observed sensitivity of the flow, and k , to small π . Furthermore, equation 7.22 is exactly consistent with the ratio formed by equation 7.17.

In summary, compared to a clean free surface, a rigid interface can reduce the gas-transfer velocity significantly through the modification of the Schmidt number dependence. This modification is a consequence of the difference in the kinematic boundary conditions between the two cases (e.g., Coantic [19]; Ledwell [82]). For O_2 , this reduction is approximately a factor of three. Based on simple scaling analysis, it is found that a surface film exhibiting surface tension gradients can change the dynamic boundary condition from a

stress-free condition to a condition much like a rigid wall for very small values of surface pressure. Thus, in terms of surface pressure, a surfactant surface can quickly resemble a rigid wall, reducing the gas transfer notably for small values of π toward a final limiting value.

7.8.5 Relationship between surface divergence and gas-transfer

Earlier, in section 2.2.2, simple scaling of the convection-diffusion equation yielded an expression for the dimensionless gas-transfer velocity in terms of the Schmidt and Reynolds numbers:

$$k^+ = \frac{k}{aL} = \mathcal{C} Re_a^{-1/2} Sc^{-n}, \quad (7.23)$$

where $Re_a = aL^2/\nu$ and the Schmidt number exponent n is governed by the surface condition. This relationship was explored using the present gas-transfer and hydrodynamic data. In order to determine the factor \mathcal{C} and the Schmidt number dependence, the transfer velocity data and the surface divergence data were fitted to the functional form $k^+ Re_a^{1/2} = \mathcal{C} Sc^{-n}$ in a least-squares sense. In this case, L was taken to be $2L_{HT}$, evaluated at the free surface. In order to characterize n in terms of a measured property of the surface, the stress ratio \mathcal{J} derived in the previous section was used. A more relevant quantity might involve the surface (visco)elasticity, which is more appropriate in describing the response of the surface to in-plane straining motions. However, as mentioned in chapter 2, such a measurement is very difficult. Since the observed effect of the presence of a surface film has been an inverse power-law, or exponential, type reduction in certain quantities of interest (e.g., figures 7-17 and 7-34), an empirical expression relating n and \mathcal{J} that captured such behavior was sought. For simplicity, the following relationship between surface pressure and n was postulated:

$$n = 2/3 - 1/6 e^{-\alpha\mathcal{J}}. \quad (7.24)$$

In this way, when $\mathcal{J} = 0$, $n = 1/2$, and the surface is clean; when $\alpha\mathcal{J} > 1$, $n \approx 2/3$ and the surface is like a rigid wall. In performing the fitting procedure, negative values of \mathcal{J} were set to zero, and values of \mathcal{J} greater than one were set to unity. The results from the fitting procedure provided the two constants α and \mathcal{C} , yielding a relationship between

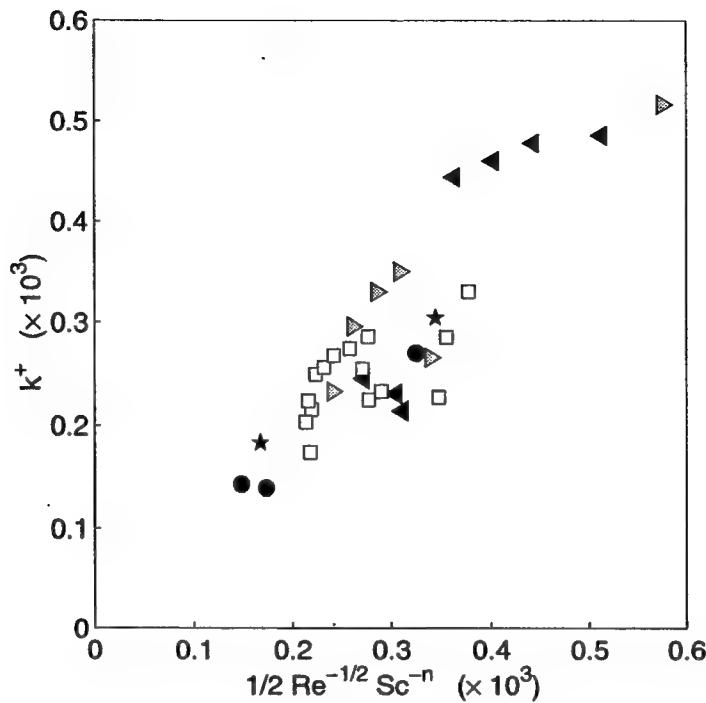


Figure 7-35: Dimensionless relationship between surface divergence and gas-transfer velocity for all data sets using equation 7.25. All markers are as in previous figures. The axis values have been scaled by 10^3 for clarity. $R^2 = 0.88$.

surface divergence and gas transfer, written as

$$k^+ = \frac{1}{2} Re_a^{-1/2} Sc^{-n}, \quad n = 2/3 - 1/6 e^{-J/2}. \quad (7.25)$$

Plotting the (k, a) data using this expression with the estimates of surface pressure as given in table 7.2 yields the result shown in figure 7-35. The relationship appears reasonably robust, with only a few data points deviating noticeably. Only the grid turbulence data are included in this figure. The R^2 statistic for this relationship was 0.88, the highest value thus far.

7.9 Summary

This chapter has looked at measurements of gas transfer across an air–water interface. By reviewing some of the earlier efforts at relating gas exchange to certain measurable quantities, it was found that there exist a number of shortcomings. Present approaches at parameterizing the gas-transfer velocity are susceptible to surface film effects and are often lacking a connecting physical mechanism. As introduced in chapter 2, the ideas of surface renewal and surface divergence were explored in this chapter. This was accomplished through contemporaneous measurements of the gas-transfer velocity and the free-surface velocity field using the surface PIV technique.

Analysis of the free-surface flow indicated that of the regions studied, there was not a great deal of variability, but that the impact of surface films was pronounced. Surfactant effects were found to dramatically alter the nature of the surface flow field, modifying the levels of turbulence, vorticity, and divergence. Significant effects were realized for small values of surface pressure ($< 1 \text{ mN/m}$), which are not uncommon in natural systems. It should be pointed out that for waterside-controlled gases (e.g., O_2) the dominant effect of surface films on the air–water gas exchange is a hydrodynamic one and is not related to the formation of a monomolecular barriers at the surface that impede interfacial transport (Liss [85]). The analysis of the coupled measurements of gas transfer and surface hydrodynamics revealed that the surface divergence appears to be a key process in air–water gas exchange, particularly in light of the results from the waves experiments. Finally, using the data presented in this chapter, a relationship between surface divergence and the gas-transfer velocity was obtained. The relationship derives from basic scaling of the convection–diffusion equation, and is novel in that it specifies the Schmidt number exponent using a dimensionless stress ratio that provides a measure of the degree to which the surface behaves as a rigid boundary.

Chapter 8

Concluding discussion

8.1 Summary of key results

This thesis has presented a new experimental approach to the challenging problem of free-surface flow measurements. The surface PIV technique developed in this work was comprised of the processing algorithms analyzed in chapter 3 and the opto-video hardware arrangement and methods described in chapter 4. Based on several tests with numerically simulated flow fields, it was shown that the calculation of the cross-correlation function used to estimate the local flow velocity with a direct spatial multiplication operation was more accurate than a calculation using a simple FFT-based correlation. However, the efficiency of the direct spatial method was 1–2 orders of magnitude worse than the FFT scheme. This prompted the development of a new hybrid correlation method, which uses an FFT-based approach to generate a first-order calculation of the correlation and then refines the result with a localized direct spatial calculation. In this way, the out-of-pattern effects that can plague a simple FFT correlation are eliminated, while improving upon the speed of the direct spatial method by roughly an order of magnitude. Tests with highly divergent simulated motions revealed that the PIV processing algorithms are capable of determining the displacement fields with very favorable accuracy.

The use of a three-dimensional light cone with the surface PIV technique is a simple, yet novel, approach to PIV flow illumination. In this way, surface particles are assured illumination, even when the free-surface is undulating as shown in chapter 5. This approach

is accompanied by concerns relating to subsurface particle illumination and the possible influence of such on the surface velocity measurement. However, by exploiting the scattering behaviors of particles on a water surface and employing proper camera filtering techniques, the contribution of illuminated particles at depth can be removed and a surface velocity measurement obtained. This methodology was validated through a series of comparative experiments that verified the accuracy of the surface PIV technique and estimated the measurement to be, at worst, in the upper $500 \mu\text{m}$, but more likely, in the upper $100 \mu\text{m}$. An attractive feature of this means of illumination emerges when wavy surfaces are studied. This is shown through the experiments with surface waves in chapter 5. For low amplitude undulations where the correction due to the surface projection is not prohibitively significant, the surface PIV technique can provide accurate measurements of the two-dimensional free-surface flow field.

A large portion of the work that contributed to this thesis centered around the grid-stirred tank that is covered in detail in chapter 6. The grid-stirred tank that was designed and constructed for this thesis work is a useful tool for turbulent mixing gas-transfer studies. In particular, because of its ability to be reasonably managed for chemical cleanliness, experiments that seek to carefully explore the role of surface films on air-water gas exchange are possible. The tank is also well-suited for velocity measurements (single point, as in ADV or LDV, or full-field with PIV) and other measurements as well. Diagnostic ADV experiments of the bulk flow showed that the grid tank generates relatively well-behaved flows for a wide range of grid forcing conditions. Initial start-up transients were found to be insignificant, the flow was observed to be reasonably repeatable between experiments, and the generated bulk turbulence was found to be stationary. Another group of bulk flow experiments with the ADV revealed that the turbulence in the grid tank (RMS turbulent velocity and integral lengthscale) was described relatively well by the empirical expressions of Hopfinger and Toly [56]. These results also indicated horizontally isotropic bulk turbulence, with only slight vertical anisotropy. Motivated by aspects of the results from the ADV measurements, the spatial variability, flow structure, and secondary flows were investigated in several vertical planes within the grid tank using conventional light sheet PIV. These measurements confirmed the existence of horizontal spatial inhomogeneity and secondary

mean motions in the grid tank. While the horizontal spatial variability was not particularly great, and the secondary circulations were typically small fractions of the local RMS turbulent velocities, claims of “zero-mean-flow” should be made carefully when dealing with grid-stirred tanks. Additional measurements with an infrared imager that looked at the nature of the grid-stirred turbulence at the free surface also indicated some amount of spatial variability and persistent flow features at the surface. These various sets of experiments provided a relatively complete picture of the characteristics of the grid-stirred turbulence and determined the approach adopted for the contemporaneous gas-transfer/surface PIV experiments of chapter 7.

A number of key results from this thesis may be found in chapter 7 where the surface PIV technique was coupled to measurements of gas exchange in the grid tank for varying degrees of surface cleanliness. It was shown that the degree of surface contamination, whether deriving from the natural content of the water (as in the case of tap water), the surface-active material accompanying untreated PIV seed particles, or through purposefully introduced surfactants, has a dramatic effect on both free-surface hydrodynamics and air-water gas transfer. The lone result from chapter 1 accompanies these observations by showing explicitly the impact of a surface film on the gas-transfer velocity. Just as important, that result also illustrated the need for measurements of the hydrodynamics at the free surface, as measurements below the interface may not reveal the critical processes involved in air-water transport of slightly soluble gases.

The results in chapter 7 add to the evidence showing the influence of surface films on interfacial hydrodynamics and air-water gas exchange. Surface turbulence, surface vorticity, and surface divergence were all affected by the presence of a surface film. Accompanying measurements of surface tension revealed that significant effects were realized for relatively small surface pressures ($\approx 1 \text{ mN/m}$). This effect, while more related to the surface viscoelasticity, results from the modification of the free-surface stress balance. The introduction of surface tension gradients, which can arise for interfacial flows with surfactants present, introduces a tangential stress at the interface that must be balanced by the viscous stress in the fluid at the interface. This new boundary condition leads to a highly dissipative surface flow that results in flow damping. Simple analysis of the stress balance led to the definition

of a dimensionless stress ratio. This ratio reveals how very small amounts of surface pressure can cause the surface to resemble a rigid boundary in a dynamic sense.

Motivated by the concept of surface renewal and gas transfer, a relationship between the gas-transfer velocity and surface divergence was explored. Results from the grid-turbulence experiments indicated a reasonably unique relationship between the transfer velocity and the surface divergence. However, an equally fair relationship was found between the surface RMS turbulent velocity and the transfer velocity. It was not until additional measurements for the case of mechanically generated surface waves were examined did the potential significance of the surface divergence emerge. These results showed that while the relationship obtained with RMS surface turbulence was an obvious improvement over that found using an estimate of the bulk flow turbulence, this relationship was dependent on the flow regime; *i.e.*, the results from the wave measurements appeared to exhibit a different $k-u$ relationship than that of the results from the grid turbulence measurements. However, the data from both the wave experiments and the grid turbulence experiments were reconciled by a single relationship between gas-transfer velocity and the 1/2-power of the surface divergence. A square-root dependence of the transfer velocity on divergence agrees with previously conceived models (*e.g.*, Csanady [22]; Ledwell [82]). Using the results of surface divergence, gas transfer, and surface pressure, a relationship between the divergence and the transfer velocity was developed. The relationship derives from basic scaling of the convection-diffusion equation, and is novel in that it specifies the Schmidt number exponent using a dimensionless stress ratio that provides a measure of the degree to which the surface behaves as a rigid boundary. This relationship is hoped to illustrate the possibly fundamental role of surface divergence in air-water gas exchange. It also is an attempt to quantify, in a physically meaningful way, the interactions between surface renewal, surfactants, and gas transfer.

8.2 Limitations and thoughts for future work

The basis for this thesis was a large body of experimental work, and as with any experimental study, there are measurement uncertainties and errors. While the surface PIV technique presented in this thesis can be a valuable tool for the study of a range of free-surface

flows, both wavy and non-wavy, the technique does have limitations. The grid turbulence experiments under cleaned surface conditions began to reveal these limits. Essentially, the technique begins to show degraded performance as the level of surface renewal, or surface bursting, becomes excessive, as is the case for cleaned water surfaces. With highly dynamic surface flows, fewer and fewer particles actually reside at the surface, and particles from below must re-supply the surface with seeding. As discussed in section 7.4.1, this results in PIV images with low particle seeding that produce velocity fields with increased noise and error. This being the case, there is a limitation to the level of surface flow energetics that can be accurately measured using the surface PIV technique used in this thesis. As the surface particle seeding falls off (*e.g.*, less than 5 particles per $32 \text{ px} \times 32 \text{ px}$ region), velocity estimates will begin to incur greater error, and although these estimates may still remain within acceptable limits, higher order measurements such as derivative quantities, spectra, correlations, and lengthscales will begin to degrade. However, such a restriction is not necessarily inhibitive, as there is much to be learned about free-surface hydrodynamics and the mechanisms of gas exchange under less energetic conditions. To be sure, gas transfer under weak forcing conditions such as in limnological environments or under low winds at sea is a topic of serious interest.

There are a number of potentially interesting avenues of study that exceeded the scope of the present investigation, but could serve as natural extensions of this work. Although, the grid tank was studied a great deal in this thesis, further investigation, as related to gas exchange, could continue. For example, it would be worthwhile to determine what effect, if any, the tank secondary circulations have on enhancing gas transfer. This same issue has arisen in annular wind-wave tanks, which develop cross-stream circulation cells due to the angular nature of the primary flow. In the case of wind-wave tanks, these effects are believed to exert little influence on gas transfer (Jähne *et al.* [64]). It also would be interesting to perform experiments with varying grid geometries in order to prescribe different integral lengthscales of the turbulence, and to explore how these scales might be related to gas transfer. Along the same lines, the question of the turbulent scales responsible for gas transfer was not examined in this investigation. In order to properly study lengthscales, and wavenumber spectra of velocity and surface divergence as well, improved PIV measurements

would be needed. In particular, use of the complete $1\text{k} \times 1\text{k}$ resolution of the camera (versus the $648 \text{ px} \times 648 \text{ px}$ that was able to be used here) would be requisite in order to collect a sufficient spatial sampling necessary for robust calculations of spectra and lengthscales. Such measurements then might allow improved comparisons between the competing ideas of Fortescue and Pearson [33] and Lamont and Scott [78]. Furthermore, measurements of integral lengthscales with and without surface films could prove insightful and examined in light of the results of Lee *et al.* [83] who report an increase in lengthscale for surfactant surfaces.

The gas-transfer aspects of this study could be extended easily with measurement of multiple gases coincidentally. This would provide an experimental determination of the Schmidt number exponent through equation 2.29. Such a measurement could be used to assess the gas-transfer relationship proposed at the conclusion of chapter 7. Related to the Schmidt number exponent measurement would be the use of different types of surfactants in this type of a study. Studying how different surfactants, both soluble and insoluble, affect the free-surface hydrodynamics and gas exchange, and the resulting variation of the Schmidt number exponent would prove useful in better understanding surfactant effects. Finally, a natural progression for this thesis would be to apply these techniques and analyses to a system more representative of ocean and lake environments: a wind-wave tank. Certainly, new challenges and difficulties would arise, but this is a necessary direction that needs to be taken in order to better simulate the processes of air-sea interaction.

Appendix A

Relating Re_ℓ and Re_λ

If Re_ℓ is known (where ℓ is a turbulent integral scale of the largest eddies), it is possible to estimate a value for Re_λ , the Reynolds number based on the longitudinal Taylor microscale. We begin with the definition

$$Re_\lambda = \frac{u_o \lambda}{\nu}, \quad (\text{A.1})$$

with $u_o^2 \equiv \bar{u^2} = \bar{v^2} = \bar{w^2}$ for isotropic, homogeneous turbulence (IHT). Use of the Von Karman–Howarth relation,

$$\lambda = \sqrt{15 \frac{\nu}{\epsilon} u_o^2}, \quad (\text{A.2})$$

where ϵ is the turbulent dissipation rate, allows us to rewrite equation A.1 as

$$Re_\lambda = \sqrt{\frac{15}{\epsilon \nu} u_o^2}. \quad (\text{A.3})$$

For IHT, $\epsilon \equiv u_o^3/\ell$. Incorporating this relation, equation A.3 becomes

$$Re_\lambda = \sqrt{\frac{15 u_o \ell}{\nu}}, \quad (\text{A.4})$$

which is simply,

$$Re_\lambda = \sqrt{15} Re_\ell^{1/2}. \quad (\text{A.5})$$

Appendix B

Uncertainty analysis for chapter 7 PIV results

In order to estimate the uncertainty in the PIV measurements of chapter 7, *i.e.*, velocity, enstrophy (vorticity), and divergence, a first-order, second-moment analysis as discussed by Taylor [128] was employed.

B.1 Velocity uncertainty

Each velocity vector measurement with PIV derives from the following expression:

$$u = \frac{\Delta p}{\Delta t} S, \quad (\text{B.1})$$

where Δp is the estimated pixel displacement of particles in the subimage, Δt is the time separation between images, and S is a length conversion factor that scales the camera pixels to physical units. If it is assumed that the uncertainties in Δp , Δt , and S are independent and random, then the fractional uncertainty in u is given by Taylor as

$$\left(\frac{\sigma_u}{|u|}\right)^2 = \left(\frac{\sigma_{\Delta p}}{|\Delta p|}\right)^2 + \left(\frac{\sigma_{\Delta t}}{|\Delta t|}\right)^2 + \left(\frac{\sigma_S}{|S|}\right)^2. \quad (\text{B.2})$$

Since the delay generator had an accuracy on the order of a nanosecond and the laser pulse timing was accurate to within ± 0.5 nanoseconds, the uncertainty in the time separation

will be assumed negligible, $\sigma_{\Delta t} \approx 0$.

The length conversion factor S was determined by imaging a scale placed in the field of view, digitizing the image, and measuring the number of pixels corresponding to a certain physical length. Therefore, the conversion is expressed as

$$S = \frac{\Delta L}{\Delta px}, \quad (\text{B.3})$$

where ΔL corresponds to the physical length measured and Δpx corresponds to the number of pixels. The fractional uncertainty in S is then

$$\left(\frac{\sigma_S}{|S|}\right)^2 = \left(\frac{\sigma_{\Delta L}}{|\Delta L|}\right)^2 + \left(\frac{\sigma_{\Delta px}}{|\Delta px|}\right)^2. \quad (\text{B.4})$$

Typically, a length $L = 5$ cm was measured, which corresponded to roughly 500 px. If it is assumed that the uncertainty in the measured physical length was $\sigma_{\Delta L} = 0.05$ cm, and the uncertainty in the measured distance in pixels was $\sigma_{\Delta px} = 1$ px, then,

$$\left(\frac{\sigma_S}{|S|}\right)^2 = 1.04 \times 10^{-4}. \quad (\text{B.5})$$

The uncertainty in the estimated pixel displacement of particles in a subimage, $\sigma_{\Delta p}$, derives from the processing of the PIV images. To assign a value to this uncertainty, aspects of chapter 3 were considered. Of the three flow types examined in chapter 3, the divergent flows in section 3.3.3 are likely to be the most representative of the flows encountered in the free-surface PIV experiments of chapter 7. The analysis of the accuracy of the hybrid correlation method (used almost exclusively in chapter 7), indicated an RMS error on the order of 0.1 px. However, these images used for the tests were populated with a relatively high particle seeding density (20 particles per 32 px \times 32 px region). Additionally, the processing was performed using 32 px \times 32 px subimages, whereas 64 px \times 64 px subimages were used most often in chapter 7. Therefore, new tests were conducted based on the estimated actual particle seeding density of the surface PIV images of chapter 7 and the actual processing parameters. Using the displacement function for flow B of section 3.3.3, additional image pairs were created that had particle densities of 2, 5, 10, and 15 particles

per $32 \text{ px} \times 32 \text{ px}$ region. Inspection of the actual surface PIV images indicated that the true particle seeding density was in the range of 5–15 particles per $32 \text{ px} \times 32 \text{ px}$ region. The RMS errors in the pixel displacements for the image pairs with 5, 10, and 15 particles per $32 \text{ px} \times 32 \text{ px}$ region were found to be 0.206 px, 0.218 px, and 0.153 px, respectively. Based on these results, the uncertainty in the estimated pixel displacement was taken to be $\sigma_{\Delta p} = 0.2 \text{ px}$.

The fractional uncertainty in the velocity can now be written as

$$\left(\frac{\sigma_u}{|u|}\right)^2 = \left(\frac{0.2}{|\Delta p|}\right)^2 + 1.04 \times 10^{-4}, \quad (\text{B.6})$$

or, as the absolute uncertainty,

$$\sigma_u = |u| \sqrt{\left(\frac{0.2}{|\Delta p_x|}\right)^2 + 1.04 \times 10^{-4}}. \quad (\text{B.7})$$

Similarly, for the orthogonal velocity component,

$$\sigma_v = |v| \sqrt{\left(\frac{0.2}{|\Delta p_y|}\right)^2 + 1.04 \times 10^{-4}}. \quad (\text{B.8})$$

As part of the data reduction in chapter 7, values of σ_u and σ_v were computed for each velocity vector in each vector field. To arrive at a single estimate of the uncertainty in the velocity, these values were spatially averaged over each vector field, and those results averaged over all 192 realizations. These final values are shown as vertical or horizontal bars in the figures of chapter 7.

B.2 Enstrophy uncertainty

By computing the fluid vorticity using equations 7.7 and 7.8, the following expression was used:

$$\begin{aligned} \omega = & \frac{1}{4\Delta x \Delta y} \left\{ \left(\frac{u_1 + u_2}{2} \right) \Delta x + \left(\frac{u_2 + u_3}{2} \right) \Delta x + \left(\frac{v_3 + v_4}{2} \right) \Delta y + \left(\frac{v_4 + v_5}{2} \right) \Delta y + \right. \\ & \left. - \left(\frac{u_5 + u_6}{2} \right) \Delta x - \left(\frac{u_6 + u_7}{2} \right) \Delta x - \left(\frac{v_7 + v_8}{2} \right) \Delta y - \left(\frac{v_8 + v_1}{2} \right) \Delta y \right\}, \end{aligned} \quad (\text{B.9})$$

where u_i and v_i , $i = 1 \dots 8$ are the velocities at the eight nodes surrounding the node to be computed, and Δx and Δy are the node spacings. This reduces to

$$\begin{aligned} \omega = & \frac{u_1 + u_2}{8\Delta y} + \frac{u_2 + u_3}{8\Delta y} + \frac{v_3 + u_4}{8\Delta x} + \frac{v_4 + v_5}{8\Delta x} + \\ & - \frac{u_5 + v_6}{8\Delta y} - \frac{u_6 + u_7}{8\Delta y} - \frac{v_7 + v_8}{8\Delta x} - \frac{v_8 + v_1}{8\Delta x}. \end{aligned} \quad (\text{B.10})$$

As a sum, the uncertainty in the vorticity can be expressed through

$$\sigma_\omega^2 = \left[\sigma \left\{ \frac{u_1 + u_2}{8\Delta y} \right\} \right]^2 + \dots - \left[\sigma \left\{ \frac{v_8 + v_1}{8\Delta x} \right\} \right]^2. \quad (\text{B.11})$$

The uncertainty in the quantity $(u_1 + u_2)/(8\Delta y)$ can be written as

$$\sigma \left\{ \frac{u_1 + u_2}{8\Delta y} \right\} = \frac{1}{8} \sigma \left\{ \frac{u_1 + u_2}{\Delta y} \right\} = \frac{1}{8} \left| \frac{u_1 + u_2}{\Delta y} \right| \sqrt{\left(\frac{\sigma_{u_1+u_2}}{u_1 + u_2} \right)^2 + \left(\frac{\sigma_{\Delta y}}{\Delta y} \right)^2}. \quad (\text{B.12})$$

Since the uncertainty in the node spacing is directly proportional to the uncertainty in the length conversion (found to be quite small), $\sigma_{\Delta y} \approx 0$ will be assumed. Equation B.12 then becomes

$$\sigma \left\{ \frac{u_1 + u_2}{8\Delta y} \right\} = \frac{1}{8} \frac{\sigma_{u_1+u_2}}{\Delta y}. \quad (\text{B.13})$$

The uncertainty in the velocity sum $\sigma_{u_1+u_2}$ is

$$\sigma_{u_1+u_2} = \sigma(u_1 + u_2) = \sqrt{\sigma_{u_1}^2 + \sigma_{u_2}^2} = \sqrt{2}\sigma_u, \quad (\text{B.14})$$

and

$$\sigma \left\{ \frac{u_1 + u_2}{8\Delta y} \right\} = \frac{\sqrt{2}}{8} \frac{\sigma_u}{\Delta y}. \quad (\text{B.15})$$

Equation B.11 is now

$$\sigma_\omega^2 = \left[\frac{\sqrt{2}}{8} \frac{\sigma_u}{\Delta y} \right]^2 + \dots - \left[\frac{\sqrt{2}}{8} \frac{\sigma_v}{\Delta x} \right]^2. \quad (\text{B.16})$$

Since $\Delta x = \Delta y$ is this case, and assuming $\sigma_u = \sigma_v$,

$$\sigma_\omega^2 = 8 \left[\frac{\sqrt{2}}{8} \frac{\sigma_u}{\Delta x} \right]^2 = \frac{1}{4} \left(\frac{\sigma_u}{\Delta x} \right)^2. \quad (\text{B.17})$$

and the absolute uncertainty in the vorticity is

$$\sigma_\omega = \frac{1}{2} \frac{\sigma_u}{\Delta x}. \quad (\text{B.18})$$

Since enstrophy is vorticity to the power 2, the fractional uncertainties in the two quantities are related as

$$\frac{\sigma_\varepsilon}{|\varepsilon|} = 2 \frac{\sigma_\omega}{|\omega|}. \quad (\text{B.19})$$

This can be rewritten for the absolute uncertainty in the enstrophy:

$$\sigma_\varepsilon = 2\sqrt{\varepsilon} \sigma_\omega. \quad (\text{B.20})$$

With equation B.18, this becomes

$$\sigma_\varepsilon = \frac{\sqrt{\varepsilon} \sigma_u}{\Delta x}. \quad (\text{B.21})$$

Using values for σ_u , the uncertainty in the enstrophy is shown in chapter 7 with vertical or horizontal bars.

B.3 Divergence uncertainty

Equation 7.9 gives the expression for the divergence as implemented in the chapter 7 analysis. For the present purposes, it will be rewritten as

$$a = \frac{u_1 - u_2}{2\Delta x} + \frac{v_1 - v_2}{2\Delta y}. \quad (\text{B.22})$$

As a sum, the uncertainty in the divergence is expressed through

$$\sigma_a^2 = \left[\sigma \left\{ \frac{u_1 - u_2}{2\Delta x} \right\} \right]^2 + \left[\sigma \left\{ \frac{v_1 - v_2}{2\Delta y} \right\} \right]^2. \quad (\text{B.23})$$

In a manner identical to that for the vorticity,

$$\sigma \left\{ \frac{u_1 - u_2}{2\Delta x} \right\} = \frac{1}{2} \frac{\sigma_{u_1 - u_2}}{\Delta x}, \quad (\text{B.24})$$

and

$$\sigma_{u_1-u_2} = \sqrt{2}\sigma_u. \quad (\text{B.25})$$

Again, since $\Delta x = \Delta y$, and assuming $\sigma_u = \sigma_v$, equation B.23 becomes

$$\sigma_a^2 = 2 \left[\frac{\sqrt{2}}{2} \frac{\sigma_u}{\Delta x} \right]^2 = \left(\frac{\sigma_u}{\Delta x} \right)^2, \quad (\text{B.26})$$

and the absolute uncertainty in the divergence is

$$\sigma_a = \frac{\sigma_u}{\Delta x}. \quad (\text{B.27})$$

To calculate the uncertainty for the square-root of the divergence, $a^{1/2}$, appearing in figure 7-28, the fractional uncertainties are used:

$$\frac{\sigma_{a^{1/2}}}{|a^{1/2}|} = \frac{1}{2} \frac{\sigma_a}{|a|}. \quad (\text{B.28})$$

This can be rewritten for the absolute uncertainty in $a^{1/2}$:

$$\sigma_{a^{1/2}} = \frac{1}{2} \sigma_a \left| \frac{1}{a^{1/2}} \right|. \quad (\text{B.29})$$

With equation B.27, this becomes

$$\sigma_{a^{1/2}} = \frac{1}{2} \frac{\sigma_u}{\Delta x} \left| \frac{1}{a^{1/2}} \right|. \quad (\text{B.30})$$

Using values for σ_u , the uncertainties in the divergence and the square-root of the divergence are shown in chapter 7 with vertical or horizontal bars.

References

- [1] N. K. Adam. *The Physics and Chemistry of Surfaces*. Oxford University Press, New York, NY, 1930.
- [2] R. J. Adrian. Multi-point optical measurements of simultaneous vectors in unsteady flow—a review. *Int. J. Heat Fluid Flow*, 7:127–145, 1986.
- [3] R. J. Adrian. Particle-imaging techniques for experimental fluid mechanics. *Ann. Rev. Fluid Mech.*, 23:261–304, 1991.
- [4] J. C. Agüí and J. Jimenez. On the performance of particle tracking. *J. Fluid Mech.*, 185:447–468, 1987.
- [5] W. Asher and R. Wanninkhof. Transient tracers and air-sea gas transfer. *J. Geophys. Res.*, 103:15939–15958, 1998.
- [6] W. E. Asher and J. F. Pankow. The interaction of mechanically generated turbulence and interfacial films with a liquid phase controlled gas/liquid transport process. *Tellus*, 38:305–318, 1986.
- [7] D. D. Back and M. J. McCready. Effect of small-wavelength waves on gas transfer across the ocean surface. *J. Geophys. Res.*, 93:5143–5152, 1988.
- [8] M. L. Banner and W. L. Peirson. An observational study of the aqueous surface layer beneath a wind-driven air-water interface. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 115–124, Heidelberg, Germany, 1995. AEON Verlag.
- [9] J. S. Bendat and A. G. Piersol. *Measurement and Analysis of Random Data*. John Wiley & Sons, Inc., New York, NY, 1966.

[10] E. J. Bock and N. M. Frew. Static and dynamic response of natural multicomponent oceanic surface films to compression and dilation: laboratory and field observations. *J. Geophys. Res.*, 98:14599–14617, 1993.

[11] E. J. Bock and T. Hara. Optical measurements of capillary-gravity wave spectra using a scanning laser slope gauge. *J. Atmos. Oceanic Technol.*, 12:395–403, 1995.

[12] H.-C. Broecker, J. Petermann, and W. Siems. The influence of wind on CO₂-exchange in a wind-wave tunnel, including the effects of monolayers. *J. Mar. Res.*, 36:595–610, 1978.

[13] B. H. Brumley and G. H. Jirka. Near-surface turbulence in a grid-stirred tank. *J. Fluid Mech.*, 183:235–263, 1987.

[14] B. H. Brumley and G. H. Jirka. Air–water transfer of slightly soluble gases: turbulence, interfacial processes and conceptual models. *PhysicoChem. Hydodyn.*, 10:295–319, 1988.

[15] P. J. Burt, C. Yen, and X. Xu. Local correlation measures for motion analysis: a comparative study. In *IEEE CVPR Conference*, pages 246–252. IEEE Compt. Soc. Press, June 1983.

[16] G. Caulliez. Measuring the wind-induced water surface flow by laser-doppler anemometry. *Exp. Fluids*, 5:145–153, 1987.

[17] W. C. Chan and L. E. Scriven. Absorption into irrotational stagnation flow: a case study in convective diffusion theory. *Ind. Engng. Chem. Fundam.*, 9:114–120, 1970.

[18] C. R. Chu and G. H. Jirka. Turbulent gas flux measurements below the air–water interface of a grid-stirred tank. *Int. J. Heat Mass Transfer*, 35:1957–1968, 1992.

[19] M. Coantic. A model of gas transfer across air/water interfaces with capillary waves. *J. Geophys. Res.*, 91:3925–3943, 1986.

[20] E. A. Cowen and S. G. Monismith. A hybrid digital particle tracking velocimetry technique. *Exp. Fluids*, 22:199–211, 1997.

[21] E. A. Cowen, S. G. Monismith, and J. R. Koseff. Digital particle tracking velocimetry measurements very near a free-surface. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 135–144, Heidelberg, Germany, 1995. AEON Verlag.

[22] G. T. Csanady. The role of breaking wavelets in air-sea gas transfer. *J. Geophys. Res.*, 95:749–759, 1990.

[23] P. V. Dankwerts. Significance of liquid-film coefficients in gas adsorption. *Ind. Eng. Chem.*, 43:1460–1467, 1951.

[24] J. T. Davies. The effects of surface films in damping eddies at a free surface of a turbulent liquid. *Proc. Roy. Soc. Lond.*, A290:515–526, 1966.

[25] J. T. Davies and E. K. Rideal. *Interfacial Phenomena*. Academic Press, New York, NY, 1963.

[26] I. P. D. De Silva and H. J. S. Fernando. Oscillating grids as a source of nearly isotropic turbulence. *Phys. Fluids*, 6:2455–2464, 1994.

[27] M. D. DeGrandpre, W. R. McGillis, N. M. Frew, and E. J. Bock. Laboratory measurements of seawater CO₂ gas fluxes. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 375–383, Heidelberg, Germany, 1995. AEON Verlag.

[28] T. D. Dickey, B. Hartman, D. Hammond, and E. Hurst. A laboratory technique for investigating the relationship between gas transfer and fluid turbulence. In W. Brutsaert and G. H. Jirka, editors, *Gas Transfer at Water Surfaces*, pages 93–100, Minneapolis, MN, 1984. D. Reidel Publishing Company.

[29] J. Dieter, F. Hering, R. Bremeyer, and B. Jähne. Measurements of velocity profiles in the aqueous boundary layer at a wind-driven water surface. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 145–152, Heidelberg, Germany, 1995. AEON Verlag.

[30] X. E and E. J. Hopfinger. On mixing across an interface in stably stratified fluid. *J. Fluid Mech.*, 166:227–244, 1986.

[31] H. J. S. Fernando and I. P. D. De Silva. Note on secondary flows in oscillating-grid, mixing box experiments. *Phys. Fluids*, 5:1849–1851, 1993.

[32] A. M. Fincham and G. R. Spedding. Low cost, high resolution DPIV measurement of turbulent fluid flow. *Exp. Fluids*, 23:449–462, 1997.

[33] G. F. Fortescue and J. R. A. Pearson. On gas adsorption into a turbulent liquid. *Chem. Eng. Sci.*, 22:1163–1176, 1967.

[34] N. M. Frew. The role of organic films in air-sea gas exchange. In P. S. Liss and R. A. Duce, editors, *The Sea Surface and Global Change*, pages 121–172. Cambridge University Press, New York, NY, 1997.

[35] N. M. Frew, E. J. Bock, W. R. McGillis, A. V. Karachintsev, T. Hara, T. Münsterer, and B. Jähne. Variation of air-water gas transfer velocity with wind stress and surface viscoelasticity. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 529–541, Heidelberg, Germany, 1995. AEON Verlag.

[36] N. M. Frew, J. C. Goldman, M. R. Dennett, and A. S. Johnson. Impact of phytoplankton-generated surfactants on air-sea gas exchange. *J. Geophys. Res.*, 95:3337–3352, 1990.

[37] N. M. Frew and R. K. Nelson. Isolation of marine microlayer film surfactants for ex situ study of their surface physical and chemical properties. *J. Geophys. Res.*, 97:5281–5290, 1992.

[38] N. M. Frew and R. K. Nelson. Scaling of marine microlayer film surface pressure-area isotherms using chemical attributes. *J. Geophys. Res.*, 97:5291–5300, 1992.

[39] Nelson M. Frew, 2000. Personal communication.

[40] G. L. Gaines. *Insoluble Monolayers at Liquid-Gas Interfaces*. John Wiley & Sons, Inc., New York, NY, 1966.

[41] J. George, F. Minel, and M. Grisenti. Physical and hydrodynamic parameters controlling gas-liquid mass transfer. *Int. J. Heat Mass Transfer*, 37:1569–1578, 1994.

[42] J. C. Goldman, M. R. Dennett, and N. M. Frew. Surfactant effects on air-sea gas exchange under turbulent conditions. *Deep Sea Res.*, 35:1953–1970, 1988.

[43] R. C. Gonzalez and P. Wintz. *Digital Image Processing*. Addison-Wesley, Reading, MA, 1987.

[44] F. C. Goodrich. On the damping of water waves by monomolecular films. *J. Phys. Chem.*, 66:1858–381, 1962.

[45] E. J. Green and D. E. Carritt. New tables for oxygen saturation of seawater. *J. Mar. Res.*, 25:140–147, 1967.

[46] J. S. Gulliver and A. Tamburrino. Turbulent surface deformation and their relationship to mass transfer in an open-channel flow. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 589–600, Heidelberg, Germany, 1995. AEON Verlag.

[47] R. A. Handler, T. F. Swean, R. I. Leighton, and J. D. Swearingen. Length scales and the energy balance for turbulence near a free surface. *AIAA J.*, 31:1998–2007, 1993.

[48] T. Hara, E. J. Bock, N. M. Frew, , and W. R. McGillis. Relationship between air-sea gas transfer velocity and surface roughness. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 611–616, Heidelberg, Germany, 1995. AEON Verlag.

[49] R. L. Hardy. Multiquadratic equations for topography and other irregular surfaces. *J. Geophys. Res.*, 76:1905–1915, 1971.

[50] H. Haussecker, S. Reinelt, and B. Jähne. Heat as a proxy tracer for gas exchange measurements in the field: principles and technical realization. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 405–413, Heidelberg, Germany, 1995. AEON Verlag.

[51] E. Hecht. *Optics*. Addison-Wesley, Reading, MA, 1987.

[52] F. Hering, C. Leue, D. Wierzimok, and B. Jähne. Particle tracking velocimetry beneath water waves. Part II: Water waves. *Exp. Fluids*, 24:10–16, 1998.

- [53] J. L. Herrin and J. C. Dutton. An investigation of LDV velocity bias correction techniques for high-speed separated flows. *Exp. Fluids*, 14:354–363, 1993.
- [54] R. Higbie. The rate of adsorption of a pure gas into a still liquid during short periods of exposure. *Trans. Am. Inst. Chem. Engrs.*, 31:365–389, 1935.
- [55] A. Hirsa, G. M. Korenowski, L. M. Logory, and C. D. Judd. Determination of surface viscosities by surfactant concentration and velocity field measurements for an insoluble monolayer. *Langmuir*, 13:3813–3822, 1997.
- [56] E. J. Hopfinger and J.-A. Toly. Spatially decaying turbulence and its relation to mixing across density interfaces. *J. Fluid Mech.*, 78:155–175, 1976.
- [57] H. Huang, D. Dabiri, and M. Gharib. On errors of digital particle image velocimetry. *Meas. Sci. Technol.*, 8:1427–1440, 1997.
- [58] H. T. Huang, H. E. Fiedler, and J. J. Wang. Limitation and improvement of PIV, part I: Limitation of conventional techniques due to deformation of particle image patterns. *Exp. Fluids*, 15:168–174, 1993.
- [59] H. T. Huang, H. E. Fiedler, and J. J. Wang. Limitation and improvement of PIV, part II: Particle image distortion, a novel technique. *Exp. Fluids*, 15:263–273, 1993.
- [60] J. C. R. Hunt. Turbulence structure and turbulent diffusion near gas-liquid interfaces. In W. Brutsaert and G. H. Jirka, editors, *Gas Transfer at Water Surfaces*, pages 67–82, Minneapolis, MN, 1984. D. Reidel Publishing Company.
- [61] J. C. R. Hunt and J. M. R. Graham. Free-stream turbulence near plane boundaries. *J. Fluid Mech.*, 84:209–235, 1978.
- [62] K. A. Hunter and P. S. Liss. Organic sea surface films. In E. K. Duursma and R. Dawson, editors, *Marine Organic Chemistry*, pages 259–298. Elsevier, New York, NY, 1981.
- [63] B. Jähne, K. O. Münnich, R. Bösinger, A. Dutzi, W. Huber, and P. Libner. On the parameters influencing air–water gas exchange. *J. Geophys. Res.*, 92:1937–1949, 1987.

[64] B. Jähne, K. O. Münnich, and U. Siegenthaler. Measurements of gas exchange and momentum transfer in a circular wind–water tunnel. *Tellus*, 31:321–329, 1979.

[65] N. G. Jerlov. *Marine Optics*. Elsevier, Amsterdam, the Netherlands, 1976.

[66] A. T. Jessup, C. J. Zappa, V. Hesany, M. R. Loewen, and M. G. Skafel. Dependence of the skin layer recovery rate on heat flux and turbulence. In B. Jähne and E. Monahan, editors, *Air-Water Gas Transfer*, pages 601–610, Heidelberg, Germany, 1995. AEON Verlag.

[67] R. D. Keane and R. J. Adrian. Theory of cross-correlation analysis of PIV images. *Appl. Sci. Res.*, 49:191–215, 1992.

[68] R. D. Keane, R. J. Adrian, and Y Zhang. Super-resolution particle imaging velocimetry. *Meas. Sci. Technol.*, 6:754–768, 1995.

[69] M. Kerker. *The Scattering of Light and Other Electromagnetic Radiation*. Academic Press, New York, NY, 1969.

[70] B.-C. Khoo and A. A. Sonin. Scalar rate correlation at a turbulent liquid free surface: a two-regime correlation for high schmidt numbers. *Int. J. Heat Mass Transfer*, 35:2233–2244, 1992.

[71] E. L. G. Kit, E. J. Strang, and H. J. S. Fernando. Measurement of turbulence near shear-free density interfaces. *J. Fluid Mech.*, 334:293–314, 1997.

[72] S. A. Kitaigorodskii. On the fluid dynamical theory of turbulent gas transfer across an air-sea interface in the presence of breaking wind-waves. *J. Phys. Oceanogr.*, 14:960–972, 1984.

[73] S. Komori, R. Nagaosa, and Y. Murakami. Turbulence structure and mass transfer across a sheared air-water interface in wind-driven turbulence. *J. Fluid Mech.*, 249:161–183, 1993.

[74] E. B. Kraus and J. A. Businger. *Atmosphere–Ocean Interaction*. Oxford University Press, New York, NY, 1994.

[75] V. N. Kudryavtsev and A. V. Soloviev. On the parameterization of the cold film on the ocean surface. *Izv. Acad. Sci. USSR Atmos. Oceanic Phys.*, 21:177–183, 1985.

[76] S. Kumar and S. Banerjee. Development and application of a hierarchical system for digital particle image velocimetry to free-surface turbulence. *Phys. Fluids*, 10:160–177, 1998.

[77] S. Kumar, R. Gupta, and S. Banerjee. An experimental investigation of the characteristics of free-surface turbulence in channel flow. *Phys. Fluids*, 10:437–456, 1998.

[78] J. C. Lamont and D. S. Scott. An eddy cell model of mass transfer into the surface of a turbulent liquid. *A. I. Ch. E. J.*, 16:513–519, 1970.

[79] P. A. Lange, B. Jähne, J. Tschiersch, and J. Ilmberger. Comparison between an amplitude-measuring wire and a slope-measuring laser water wave gauge. *Rev. Sci. Instrum.*, 53:651–655, 1982.

[80] G. S. Lapham, D. R. Dowling, and W. W. Schultz. In situ force-balance tensiometry. *Exp. Fluids*, 27:157–166, 1999.

[81] C. N. S. Law, B. C. Khoo, and T. C. Chew. Turbulence structure in the immediate vicinity of the shear-free air–water interface induced by a deeply submerged jet. *Exp. Fluids*, 27:321–331, 1999.

[82] J. R. Ledwell. The variation of the gas transfer coefficient with molecular diffusivity. In W. Brutsaert and G. H. Jirka, editors, *Gas Transfer at Water Surfaces*, pages 293–302, Minneapolis, MN, 1984. D. Reidel Publishing Company.

[83] Y. H. Lee, G. T. Tsao, and P. C. Wankat. Hydrodynamic effect of surfactants on gas-liquid oxygen transfer. *Am. Inst. Chem. Engrs. J.*, 26:1008–1012, 1980.

[84] J. C. Lin and D. Rockwell. Evolution of a quasi-steady breaking wave. *J. Fluid Mech.*, 302:29–44, 1995.

[85] P. S. Liss. Gas transfer, experiments and geochemical applications. In P. S. Liss and G. N. Slinn, editors, *Air-Sea Exchange of Gases and Particles*. Reidel, Higham, MA, 1983.

[86] P. S. Liss and L. Merlivat. Air-sea gas exchange rates: Introduction and synthesis. In P. Buat-Menard, editor, *The Role of Air-Sea Exchange in Geochemical Cycling*, pages 113–127. D. Reidel Publishing Company, Norwell, MA, 1986.

[87] P. S. Liss and P. G. Slater. Flux of gases across the air-sea interface. *Nature*, 247:181–184, 1974.

[88] P. M. Llyod, P. K. Stansby, and D. J. Ball. Unsteady surface-velocity field measurement using particle tracking velocimetry. *J. Hydraul. Res.*, 33:519–534, 1995.

[89] L. M. Logory, A. Hirsa, and D. G. Anthony. Interaction of wake turbulence with a free surface. *Phys. Fluids*, 8:805–815, 1996.

[90] E. H. Lucassen-Reynders and J. Lucassen. Properties of capillary waves. *Advan. Colloid Interface Sci.*, 2:347–395, 1969.

[91] J. T. Mass. Dynamic properties of seawater surfactants. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 1995.

[92] M. J. McCready, E. Vassiliadou, and T. J. Hanratty. Computer simulation of turbulent mass transfer at a mobile interface. *Am. Inst. Chem. Engrs. J.*, 32:1108–1115, 1986.

[93] T. J. McDougall. Measurements of turbulence in a zero-mean-shear mixed layer. *J. Fluid Mech.*, 94:409–431, 1979.

[94] W. R. McGillis, J. W. H. Dacey, N. M. Frew, E. J. Bock, and R. K. Nelson. Water-air flux of dimethylsulfide. *J. Geophys. Res.*, 105:1187–1193, 2000.

[95] Wade R. McGillis, 1999. Personal communication.

[96] S. P. McKenna. The influence of surface films on interfacial flow dynamics. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 1997.

[97] W. McKeown and R. Leighton. Mapping heat flux. *J. Atmos. Oceanic Technol.*, 16:80–91, 1999.

[98] D. K. McLaughlin and W. G. Tiederman. Biasing correction for individual realization of laser anemometer measurements in turbulent flow. *Phys. Fluids*, 16:2082–2088, 1973.

[99] W. McLeish and G. E. Putland. Measurements of wind-driven flow profiles in the top millimeter of water. *J. Phys. Oceanogr.*, 5:516–518, 1975.

[100] W. K. Melville and R. J. Rapp. The surface velocity field in steep and breaking waves. *J. Fluid Mech.*, 189:4–22, 1988.

[101] J. H. Milgram. Short wave damping in the simultaneous presence of a surface film and turbulence. *J. Geophys. Res.*, 103:15717–15727, 1998.

[102] D. B. Moog and G. H. Jirka. Air-water gas transfer in uniform channel flow. *J. Hydraul. Eng.*, 125:3–10, 1999.

[103] Robert K. Nelson, 1996. Personal communication.

[104] K. Okamoto, Y. A. Hassan, and W. D. Schmidl. New tracking algorithm for particle image velocimetry. *Exp. Fluids*, 19:342–347, 1995.

[105] K. Okuda, S. Kawai, M. Tokuda, and Y. Toba. Detailed observation of the wind-exerted surface flow by use of flow visualization methods. *J. Oceanogr. Soc. Japan*, 32:53–64, 1976.

[106] W. L. Peirson. Measurement of surface velocities and shears at a wavy air–water interface using particle image velocimetry. *Exp. Fluids*, 23:427–437, 1997.

[107] G. M. Quénod, J. Pakleza, and T. A. Kowalewski. Particle image velocimetry with optical flow. *Exp. Fluids*, 25:177–189, 1998.

[108] M. Raffel, C. Willert, and J. Kompenhans. *Particle Image Velocimetry: A Practical Guide*. Springer-Verlag, Berlin, Germany, 1998.

[109] M. Rashidi. Burst-interface interactions in free surface turbulent flows. *Phys. Fluids*, 9:3485–3501, 1997.

[110] M. Rashidi, G. Hetsroni, and S. Banerjee. Meachanisms of heat and mass transport at gas-liquid interfaces. *Int. J. Heat Mass Transfer*, 34:1799–1810, 1991.

[111] J. E. Rehm and N. T. Clemens. An improved method for enhancing the resolution of conventional double-exposure single-frame particle image velocimetry. *Exp. Fluids*, 26:497–504, 1999.

[112] T. Roesgen and R. Totaro. Two-dimensional on-line particle imaging velocimetry. *Exp. Fluids*, 19:188–193, 1995.

[113] H. Rouse and J. Dodu. Diffusion turbulente à travers une discontinuité de densité. *Houille Blanche*, 10:522–532, 1955.

[114] R. H. Sabersky, A. J. Acosta, and E. G. Hauptmann. *Fluid Flow: A First Course in Fluid Mechanics*. Macmillan Publishing Company, New York, NY, 1989.

[115] J. R. Saylor and R. A. Handler. Gas transport across an air/water interface populated with capillary waves. *Phys. Fluids*, 9:2529–2541, 1997.

[116] F. Scarano and M. L. Riethmuller. Iterative multigrid approach in PIV image processing with discrete offset. *Exp. Fluids*, 26:513–523, 1999.

[117] H. Schlichting. *Boundary Layer Theory*. McGraw-Hill, Inc., New York, NY, 1960.

[118] J. C. Scott. The preparation of clean water surfaces for fluid mechanics. In K. L. Mittal, editor, *Surface Contamination: Genesis, Detection, and Control I*, pages 477–497. Plenum Press, New York, NY, 1979.

[119] L. Shen, G. S. Triantafyllou, and D. K. P. Yue. Turbulent diffusion near a free surface. *J. Fluid Mech.*, 407:145–166, 2000.

[120] L. Shen, X. Zhang, D. K. P. Yue, and G. S. Triantafyllou. The surface layer for free-surface turbulent flows. *J. Fluid Mech.*, 386:167–212, 1999.

[121] K. S. Shifrin. *Physical Optics of Ocean Water*. American Institute of Physics, New York, NY, 1988.

[122] I. P. D. De Silva and H. J. S. Fernando. Some aspects of mixing in a stratified turbulent patch. *J. Fluid Mech.*, 240:601–625, 1992.

[123] A. V. Soloviev and P. Schlüssel. Parameterization of the cool skin of the ocean and of the air–ocean gas transfer on the basis of modeling surface renewal. *J. Phys. Oceanogr.*, 24:1339–1346, 1994.

[124] X. Song, F. Yamamoto, M. Iguchi, and Y. Murai. A new tracking algorithm of PIV and removal of spurious vectors using Delaunay tessellation. *Exp. Fluids*, 26:371–380, 1999.

[125] T. G. Springer and R. L. Pickford. Influence of surface turbulence and surfactants on gas transport through liquid interfaces. *Ind. Engng. Chem. Fundam.*, 9:458–465, 1970.

[126] N. M. Suoja. *Directional wavenumber characteristics of short sea waves*. PhD dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1999.

[127] A. Tamburrino and J. S. Gulliver. Free-surface turbulence measurements in an open-channel flow. In E. P. Rood and J. Katz, editors, *Free-Surface Turbulence*, pages 103–112, Lake Tahoe, NV, 1994. ASME.

[128] J. R. Taylor. *An Introduction to Error Analysis*. University Science Books, Mill Valley, CA, 1982.

[129] H. Tennekes and J. L. Lumley. *A First Course in Turbulence*. The MIT Press, Cambridge, MA, 1972.

[130] S. M. Thompson and J. S. Turner. Mixing across an interface due to turbulence generated by an oscillating grid. *J. Fluid Mech.*, 67:349–368, 1975.

[131] W. Thompson (Lord Kelvin). *Philos. Mag. Ser. 4*, 41:362, 1871.

[132] W.-T. Tsai. Impact of a surfactant on a turbulent shear layer under the air-sea interface. *J. Geophys. Res.*, 101:28557–28568, 1996.

- [133] W. T. Tsai and D. K. P. Yue. Effects of soluble and insoluble surfactant on laminar interactions of vortical flows with a free surface. *J. Fluid Mech.*, 289:315–349, 1995.
- [134] R. J. Volino and G. B. Smith. Use of simultaneous IR temperature measurements and DPIV to investigate thermal plumes in a thick layer cooled from above. *Exp. Fluids*, 27:70–78, 1999.
- [135] J. Voulgaris and J. H. Trowbridge. Evaluation of the acoustic Doppler velocimeter (ADV) for turbulence measurements. *J. Atmos. Oceanic Technol.*, 15:272–289, 1998.
- [136] R. Wanninkhof. Relationship between wind speed and gas exchange over the ocean. *J. Geophys. Res.*, 97:7373–7382, 1992.
- [137] R. Wanninkhof and W. R. McGillis. A cubic relationship between air-sea CO₂ exchange and wind speed. *Geophys. Res. Lett.*, 26:1889–1892, 1999.
- [138] A. Weigand. Simultaneous mapping of the velocity and deformation field at a free surface. *Exp. Fluids*, 20:358–364, 1996.
- [139] J. Westerweel. *Digital Particle Image Velocimetry: Theory and Application*. Delft University Press, Delft, the Netherlands, 1993.
- [140] J. Westerweel. Fundamentals of digital particle image velocimetry. *Meas. Sci. Technol.*, 8:1379–1392, 1997.
- [141] J. Westerweel, D. Dabiri, and M. Gharib. The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings. *Exp. Fluids*, 23:20–28, 1997.
- [142] F. M. White. *Viscous Fluid Flow*. McGraw-Hill, Inc., New York, NY, 1991.
- [143] W. G. Whitman. The two-film theory of gas adsorption. *Chem. Metall. Engrng.*, 29:146–148, 1923.
- [144] C. E. Willert and M. Gharib. Digital particle image velocimetry. *Exp. Fluids*, 10:181–193, 1991.

- [145] C. J. Zappa. *Microscale wave breaking and its effect on air-water gas transfer using infrared imagery*. PhD dissertation, University of Washington, Seattle, WA, 1999.
- [146] C. J. Zappa, A. T. Jessup, and H. Yeh. Skin layer recovery of free-surface wakes: Relationship to surface renewal and dependence on heat flux and background turbulence. *J. Geophys. Res.*, 102:21711–21722, 1998.

DOCUMENT LIBRARY

Distribution List for Technical Report Exchange – July 1998

University of California, San Diego
SIO Library 0175C
9500 Gilman Drive
La Jolla, CA 92093-0175

Hancock Library of Biology & Oceanography
Alan Hancock Laboratory
University of Southern California
University Park
Los Angeles, CA 90089-0371

Gifts & Exchanges
Library
Bedford Institute of Oceanography
P.O. Box 1006
Dartmouth, NS, B2Y 4A2, CANADA

NOAA/EDIS Miami Library Center
4301 Rickenbacker Causeway
Miami, FL 33149

Research Library
U.S. Army Corps of Engineers
Waterways Experiment Station
3909 Halls Ferry Road
Vicksburg, MS 39180-6199

Marine Resources Information Center
Building E38-320
MIT
Cambridge, MA 02139

Library
Lamont-Doherty Geological Observatory
Columbia University
Palisades, NY 10964

Library
Serials Department
Oregon State University
Corvallis, OR 97331

Pell Marine Science Library
University of Rhode Island
Narragansett Bay Campus
Narragansett, RI 02882

Working Collection
Texas A&M University
Dept. of Oceanography
College Station, TX 77843

Fisheries-Oceanography Library
151 Oceanography Teaching Bldg.
University of Washington
Seattle, WA 98195

Library
R.S.M.A.S.
University of Miami
4600 Rickenbacker Causeway
Miami, FL 33149

Maury Oceanographic Library
Naval Oceanographic Office
Building 1003 South
1002 Balch Blvd.
Stennis Space Center, MS, 39522-5001

Library
Institute of Ocean Sciences
P.O. Box 6000
Sidney, B.C. V8L 4B2
CANADA

National Oceanographic Library
Southampton Oceanography Centre
European Way
Southampton SO14 3ZH
UK

The Librarian
CSIRO Marine Laboratories
G.P.O. Box 1538
Hobart, Tasmania
AUSTRALIA 7001

Library
Proudman Oceanographic Laboratory
Bidston Observatory
Birkenhead
Merseyside L43 7 RA
UNITED KINGDOM

IFREMER
Centre de Brest
Service Documentation - Publications
BP 70 29280 PLOUZANE
FRANCE

REPORT DOCUMENTATION PAGE		1. REPORT NO. MIT/WHOI 00-29	2.	3. Recipient's Accession No.				
4. Title and Subtitle Free-Surface Turbulence and Air-Water Gas Exchange			5. Report Date September 2000	6.				
7. Author(s) Sean Patrick McKenna			8. Performing Organization Rept. No.					
9. Performing Organization Name and Address MIT/WHOI Joint Program in Oceanography/Applied Ocean Science & Engineering			10. Project/Task/Work Unit No. MIT/WHOI 00-29	11. Contract(C) or Grant(G) No.				
			(C)	(G)				
12. Sponsoring Organization Name and Address WHOI Ocean Ventures Fund			13. Type of Report & Period Covered Ph.D. Thesis	14.				
15. Supplementary Notes This thesis should be cited as: Sean Patrick McKenna, 2000. Free-Surface Turbulence and Air-Water Gas Exchange. Ph.D. Thesis. MIT/WHOI, 00-29.								
16. Abstract (Limit: 200 words) This thesis investigates the physical mechanisms of air-water gas transfer through direct measurements of turbulence at the air-water interface. To enable this study, a new approach to the particle image velocimetry (PIV) technique is developed in order to quantify free-surface flows. PIV measurements of free-surface turbulence are performed with concurrent measurements of gas transfer in a specially constructed oscillating grid-stirred tank for a range of turbulent mixing and surface conditions. Surface turbulence, vorticity, divergence, and gas transfer are all affected by the presence of a surface film, with significant effects realized for relatively small surface pressures. Results show that while a relationship between surface turbulence and the gas-transfer velocity is an obvious improvement over that found using an estimate of the bulk flow turbulence, this relationship is dependent on the flow regime. However, the data from the grid turbulence experiments, and additional data from experiments with waves, can be reconciled by a single relationship between the gas-transfer velocity and the 1/2-power of the surface divergence, which agrees with previous conceptual models. These results demonstrate the important role of surface divergence in air-water gas exchange, and relate, in a physically meaningful way, the interactions between surface renewal, surfactants, and gas transfer.								
17. Document Analysis								
<table border="0"> <tr> <td>a. Descriptors</td> </tr> <tr> <td>Turbulence</td> </tr> <tr> <td>free surface</td> </tr> <tr> <td>transport</td> </tr> </table>					a. Descriptors	Turbulence	free surface	transport
a. Descriptors								
Turbulence								
free surface								
transport								
<table border="0"> <tr> <td>b. Identifiers/Open-Ended Terms</td> </tr> </table>					b. Identifiers/Open-Ended Terms			
b. Identifiers/Open-Ended Terms								
<table border="0"> <tr> <td>c. COSATI Field/Group</td> </tr> </table>					c. COSATI Field/Group			
c. COSATI Field/Group								
18. Availability Statement Approved for publication; distribution unlimited.		19. Security Class (This Report) UNCLASSIFIED	21. No. of Pages 317					
		20. Security Class (This Page)	22. Price					